Level-1 Jet Trigger Studies for the CMS Experiment

James Brooke

A thesis submitted to the University of Bristol in accordance with the requirements of the degree of PhD in the Faculty of Science.

~36,000 words

Department of Physics

September 2002

Abstract

The Compact Muon Solenoid (CMS) detector is introduced, with particular emphasis on the calorimeters, and the trigger and data acquisition system. An FPGA-based sort processor for use in the CMS Global Calorimeter Trigger has been designed. The algorithm used and its implementation are described, together with results from a demonstrator board built to test the design. Further successful results from a second, more sophisticated prototype processor board are also described.

The Level-1 jet trigger rate and performance have been calculated using detailed simulation programs. The results are presented for low LHC luminosity running conditions. The trigger segmentation of the very forward calorimeters has been investigated. The results show that a proposed extension of the baseline segmentation (increasing the number of towers in pseudorapidity from four to six), while offering slightly improved performance, does not provide sufficient increase to warrant the change. Finally, a simple di-jet trigger can be extended using cuts on the separation of the two jets in pseudorapidity. The performance of such a trigger in selecting weak boson fusion events is found to be insufficient to be of general use. However, excellent trigger efficiency for an invisibly decaying light Higgs boson can be provided by requiring two tag jets, well separated in pseudorapidity, together with missing transverse energy above a threshold of 60 GeV.

Acknowledgements

I would like to express my gratitude to my supervisor, Greg Heath, for his guidance during my time as a graduate student, and for his excellent proof-reading of this thesis. Thanks are also due to the other members of the GCT team, in particular Dave Newbold and David Cussans, for providing ideas and assistance throughout the course of my study. I would also like to thank the members of the CMS JetMEt group, in particular Sarah Eno, Sasha Nikitenko and Sridhara Dasu for their advice and ideas on the subjects of jets and missing energy. There are many others who deserve a mention, including the RAL GCT team, the CMS monte-carlo production team, and members of the CMS software group.

I would not have finished this thesis were it not for the unending support I have received from my family and friends, for which I am eternally grateful. This thesis is dedicated to you all; thank-you!

Author’s Declaration

I declare that the work in this thesis was carried out in accordance with the Regulations of the University of Bristol. The work is original except where indicated by special reference in the text and no part of the thesis has been submitted for any other degree.

Any views expressed in this thesis are those of the author and in no way represent those of the University of Bristol. The thesis has not been presented to any other University for examination, either in the United Kingdom or overseas.

Signed

Date

Table of Contents

List of Figures

List of Tables

Chapter 1 - Introduction

This thesis contains details of several pieces of work undertaken by the author as a member of the Compact Muon Solenoid (CMS) collaboration. The CMS collaboration comprises more than 1800 scientists, from 153 institutes in 33 countries. Their common goal is the design, construction and operation of a general purpose detector at the Large Hadron Collider (LHC).

The LHC is a proton-proton collider that is currently under construction at CERN. It will explore higher energies than have previously been available to particle physicists. The first collisions are scheduled for April 2007. The CMS detector is a general purpose detector that will study both proton-proton physics and heavy ion collisions. Three other experiments are planned for the LHC: another general purpose detector for *p-p* physics (ATLAS); a specialised detector for studying b-physics (LHC-b); and a specialised detector for heavy ion collisions (ALICE).

1.1 Motivation

An overview of the Standard Model of particle physics will be given in Chapter 2, together with one of the more dominant theories that extends physics beyond the Standard Model. Leaving the matter of neutrino mass aside for the moment, the experimental evidence available to date shows excellent agreement with the Standard Model. The results of data collected at LEP, Tevatron, HERA, PEP-II and many smaller low energy experiments are found to fit the Standard Model well. The recent precision measurement of the anomalous magnetic moment of the muon by the E281 collaboration may indicate a departure from the Standard Model prediction, but this result does not yet show a statistically significant deviation. The neutrino oscillation results from SuperKamiokande and the Sudbury Neutrino Observatory do indicate a significant departure from the Standard Model as it stood a few years ago, but it seems that neutrino masses can be accomodated in the Standard Model in a consistent way without assuming new particles or interactions.

1.1.1 Searching for the Higgs

Despite the excellent agreement of experimental evidence with the Standard Model, it remains that one of its fundamental particles, the Higgs boson, has not yet been discovered. The search for the Higgs is the primary goal of particle physics today. Precision measurements made at LEP indicate that a Higgs, if it exists, will be found with mass below 200 GeV. It had been hoped that the final year of running at LEP would provide a glimpse of

the Higgs, but the final analysis eventually resulted in a lower bound on the Higgs mass, with no statistically significant signal. The search for the Higgs is now being carried out at Run-II of the Tevatron. If this proves unsuccessful, the LHC will extend the search over the full mass range. It is safe to say that if the Higgs exists, it will be discovered by the end of the LHC's lifetime.

1.1.2 Beyond the Standard Model

Although the primary aim of the LHC is to prove or disprove the existence of the Higgs, it is likely to a whole wealth of physics will become available at the LHC. Supersymmetry is perhaps the most popular extension of the Standard Model, and the LHC will be able to cover a wide range of its parameter space. Other well motivated theories, such a those of large extra dimensions, may also be within reach of the LHC.

1.1.3 An Invisible Higgs

The most popular supersymmetric models (those that preserve R-parity) imply the existence of a lightest supersymmetric particle. The other SUSY particle are unstable, and will eventually decay to the LSP (usually via a chain of other unstable particles). R-parity is favoured by limits on the proton lifetime, and because the LSP is a viable candidate for dark matter, which is required to explain various astrophysical measurements. Since the LSP does not interact with particle detectors, its presence in an LHC collision must be inferred through the resulting imbalance of energy flow, known as 'missing energy'.

If the Higgs boson is sufficiently massive, it may itself decay into such particles. If the LSP is indeed kinematically available, this is thought to be the dominant decay mode of a SUSY Higgs. Furthermore, the neutrino mass measurements have opened up other extensions to the Standard Model that can result in invisible decays of the Higgs. There is therefore a strong need to ensure such an invisibly decaying Higgs will be observable at the LHC.

1.1.4 Triggers

The studies described in this thesis are centred around the CMS trigger system. The trigger is central to the success of the experiment, as it provides the first stage of data analysis and makes the decision whether to store each event for offline analysis. Since the trigger decision may result in the event under consideration being permanently rejected, the trigger must capture important physics signals with excellent efficiency. Given that the extremely low cross-section for these signals can result in as little as a few events per year, the demands placed on the trigger are high. Furthermore, the LHC represents a very difficult environment

for the trigger, due to the presence of multiple proton-proton events in each bunch crossing. Ensuring the trigger is capable of fulfilling its role is therefore an essential ingredient of a successful experiment.

1.2 Definitions

Events & crossings

When discussing the simulation of LHC physics, a single *pp* interaction is referred to as an **event**, while a *bunch-bunch* interaction, which is generally the sum of several *pp* interactions, is referred to as a **crossing**.

Coordinates

Two co-ordinate systems are used in this thesis. The physical implementation of the experiment is usually dscribed in terms of a standard Cartesian system, where the z-direction is aligned with the beampipe, and the y-direction is vertical. When discussing the physics, the (η, ϕ, z) co-ordinates are used, where z is the same as the Cartesian system, ϕ is the azimuthal angle, given by $\phi = \text{atan}\frac{y}{x}$, and η is the *pseudorapidity*, defined by

$$\eta = -\log\left(\tan\frac{\theta}{2}\right) \tag{1.1}$$

where θ is the polar angle, given by $\theta = \text{atan}\frac{y}{z}$.

Notation

The symbol $\oplus$ is used throughout to indicate addition in quadrature.

Units

When discussing the physics of fundamental particles and their detection, energy is measured in electron volts, *eV*, and natural units are used for the other quantities. The natural units are defined by setting

$$\frac{h}{2\pi} = c = 1 \tag{1.2}$$

where h is Planck's constant and c is the speed of light in the vacuum. The units of momentum and time are therefore eV and $(eV)^{-1}$ respectively.

When discussing the implementation of the CMS detector and associated equipment, SI units are used.

Chapter 2 - Theoretical Background

This first section of this chapter gives a brief outline of the Standard Model, upon which modern particle physics is based. This is followed by a description of the problem of electroweak symmetry breaking. Discovering the nature of this mechanism is perhaps the most important challenge to particle physics today.

The Standard Model is viewed as inadequate for a variety of reasons, and the theoretical arguments that it is an approximation to a more fundamental theory are strong. The most popular theories that resolve some of these issues involve the introduction of a new symmetry between fermions and bosons, known as Supersymmetry (SUSY). An overview of SUSY is given in §2.3.

Finally, certain regions of the SUSY parameter space give rise to invisible decays of the Higgs boson, as can various other extensions to the Standard Model. The mechanisms involved are described in §2.4.

2.1 The Standard Model

The standard model of particle physics describes the fundamental particles of nature, and the interactions between them, using quantum field theory. The standard model includes all known fundamental forces with the exception of gravity. The particles of the standard model fall into two categories; the spin-1/2 fermions are often thought of as ‘matter’ particles, while the integral spin bosons can be thought of as the ‘force’ particles. The Standard Model particles and their properties are given in Table 2.1.

The fermions are subdivided into 2 groups; the leptons and the quarks. The essential difference between the two groups is that the quarks carry colour charge and therefore feel the strong force, while the leptons do not. While the leptons can exist in free space, the nature of the strong force ensures that quarks are only observed in bound states of 2 or 3 quarks, called hadrons. The quarks and leptons both come in three ‘families’ or ‘generations’, distinguishable from each other by their mass.

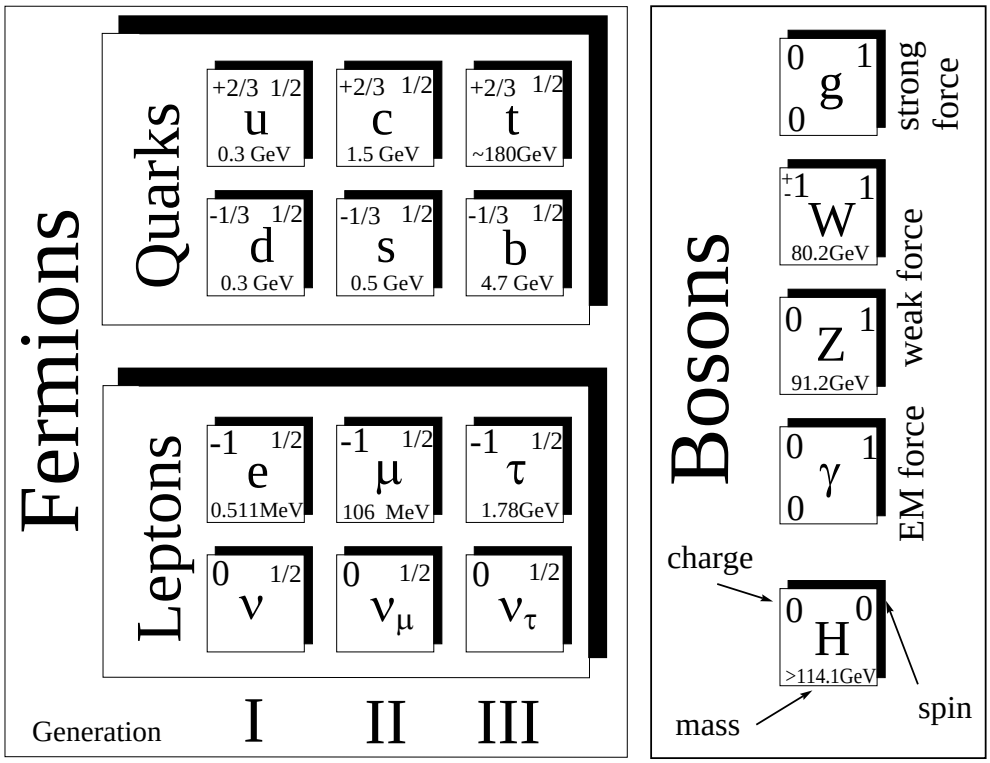

Table 2.1 - The fundamental particles of the Standard Model. Experimental limits on the neutrino masses have been omitted.

The electromagnetic force mediates interactions between charged particles. It is completely described by the theory of quantum electrodynamics (QED), a field theory with local $U(1)$ invariance. This local (or gauge) invariance requires the existence of one or more bosons that carry the force between the interacting particles. In the case of electromagnetism, the single gauge boson is the photon.

In the standard model, the weak nuclear force is unified with electromagnetism in the Glashow-Weinberg-Salam (GWS) theory of electroweak interactions. This theory contains local $SU(2) \otimes U(1)$ gauge invariance, which requires 4 gauge fields. The physical particles that arise from the mixing of these fields are the photon, the $W^{\pm}$ and the Z^0 bosons.

The strong force is described by quantum chromodynamics, a theory with local SU(3) gauge invariance. The gauge bosons are the 8 gluons, which carry the force between 'colour' charges. The gauge bosons of QCD feel the force themselves, as they are colour charged. It is this detail that results in quark confinement and the formation of hadrons.

The final particle of the Standard Model is the Higgs boson. Although the existence of this

particle has not been confirmed, it is included in the model primarily to explain how the $W^{\pm}$ and Z^0 bosons may acquire mass without breaking the gauge invariance required by the electroweak theory. Its presence can also be used to explain the fermion masses. The Higgs mechanism will be explained in more detail in §2.2.1.

2.2 Electroweak Symmetry Breaking

At the heart of the Standard Model is the GWS theory of electromagnetic and weak forces. This theory is based on a $SU(2) \otimes U(1)$ symmetry of weak isospin and hypercharge. In order to maintain local symmetry, four vector fields are introduced. The physical photon, $W^{\pm}$ and Z bosons are found to be mixtures of these gauge fields. The explicit introduction of mass terms for the physical fields into the Lagrangian renders the theory non-renormalizable. However, if *spontaneous* breaking of the $SU(2) \otimes U(1)$ symmetry is required these fields may acquire mass in a renormalizable manner.

2.2.1 The Higgs Mechanism

In order to demonstrate how a gauge symmetry may be spontaneously broken, and how this leads to massive gauge bosons, a $U(1)$ theory for a complex scalar field Φ is taken as an example. If a gauge field, A_μ, is included the theory will maintain local $U(1)$ symmetry. This is essentially QED for charged scalars, and the Lagrangian is

$$L = (\partial_\mu - igA_\mu)\Phi^*(\partial^\mu + igA^\mu)\Phi - -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - V(\Phi^*\Phi) \tag{2.1}$$

where the potential V is given by

$$V(\Phi) = \lambda|\Phi^*\Phi|^2 - \mu^2\Phi^*\Phi \tag{2.2}$$

Note that this potential has a minimum at

$$\Phi = e^{i\theta}\frac{v}{\sqrt{2}} \tag{2.3}$$

where $v = \mu/\sqrt{\lambda}$. The vacuum is now degenerate, since θ may take any value between 0 and 2π for each point in space-time. In choosing the true physical vacuum from the infinite set of possible states, the $U(1)$ symmetry is broken. The equivalence of fixing the gauge (which is necessary in order to fully quantise the theory) and choosing the physical vacuum should be noted.

Since our choice of physical vacuum is arbitrary, given $\Phi(x)$ we may always choose $\theta(x)$ for which Φ is real (this is known as the unitary gauge). We may then expand the symmetry broken field, Φ', around the physical vacuum as

$$\Phi' = \frac{1}{\sqrt{2}}(v + H) \tag{2.4}$$

where H is a real scalar field. Hence, after symmetry breaking, Eq. 2.2 becomes

$$V = \lambda v^2 H^2 + \lambda v H^3 + \frac{\lambda}{4}H^4 - \frac{\lambda v^4}{4} \tag{2.5}$$

If we now split the Lagrangian into free and interaction parts, we obtain

$$L_{free} = \frac{1}{2}\partial_\mu H \partial^\mu H - \lambda v^2 H^2 + \frac{1}{2}g^2 v^2 A_\mu A^\mu - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \tag{2.6}$$

$$L_{int} = g^2 v A_\mu A^\mu H + \frac{1}{2}g^2 A_\mu A^\mu H^2 - \lambda v H^3 - \frac{\lambda}{4}H^4 \tag{2.7}$$

We interpret this as a scalar field H of mass λv^2, and a vector field A with mass $g^2 v^2 / 2$. The interactions between the two fields are shown in Figure 2.1.

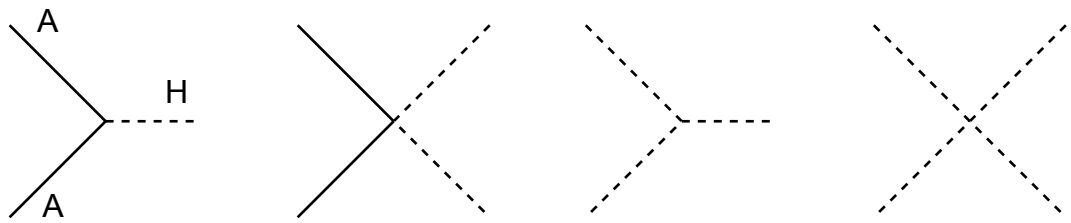

Figure 2.1 - Interactions between the vector (A) and scalar (H) fields described in §2.2.1.

So, despite the absence of an explicit mass term for the gauge field in the Lagrangian, we have obtained a massive physical particle by introducing a complex scalar field. The mechanism described above is known as the Higgs mechanism after its inventor [1]. The complex scalar field results in a single physical scalar, known as the Higgs boson, which interacts with the gauge field and with itself. It has been shown that the theory is fully renormalisable [2], which would not have been the case if the gauge boson mass had been explicitly introduced into the Lagrangian.

2.2.2 The Electroweak Model for Leptons

For the sake of simplicity, the quarks will be ignored for the remainder of §2.2, along with the 2nd and 3rd generations of leptons. The theory is easily extended to include these particles.

The electron and its neutrino are represented by a weak isospin doublet containing the left-handed particles, together with a weak isospin singlet containing the right-handed electron.

$$\psi_L = \begin{bmatrix} \nu_{eL} \\ e_L \end{bmatrix}, \psi_R = e_R \tag{2.8}$$

If a right-handed neutrino exists, it is believed to be sterile and will be ignored until §2.2.6. It is noted that an explicit mass term for fermions must mix left and right-handed states; therefore the two may only be treated separately in the massless limit. The fermion masses will instead be explained in terms of their Yukawa couplings to the Higgs field (see §2.2.4).

In order to include electromagnetic interactions in the theory, the theory must contain an unbroken $U(1)$ gauge symmetry. However, the underlying symmetry is not directly associated with electric charge. Instead it is associated with the hypercharge, Y, which is defined by

$$Q = T_3 + Y \tag{2.9}$$

where Q is electric charge and T_3 is the third component of isospin. The left-handed electron and neutrino have hypercharge -½, while the right-handed electron (having isospin of 0) has hypercharge -1.

The generators of $SU(2) \otimes U(1)$ are

$$T^1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, T^2 = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, T^3 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, Y = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{2.10}$$

Now, four gauge fields are required in order to maintain local symmetry, one for each generator. The covariant derivative for the left-handed fields is

$$D_\mu = (\partial_\mu + igW_\mu^\alpha T^\alpha + ig'B_\mu Y) \tag{2.11}$$

where the gauge fields are W_μ^α and B_μ. Since it has zero isospin, the right-handed electron has the following covariant derivative

$$D_\mu = (\partial_\mu + ig'B_\mu I) \tag{2.12}$$

The relative strength of the isospin and hypercharge interactions is characterised by the Weinberg angle, θ_W.

$$g' = g\tan\theta_W \tag{2.13}$$

The next section will explain how the massive W and Z fields arise after spontaneous breaking of the $SU(2)_L \otimes U(1)_Y$ symmetry, leaving an unbroken $U(1)_{EM}$ symmetry with a massless photon.

2.2.3 The W and Z Bosons

The charged vector bosons can be identified immediately, as it is possible to rewrite the gauge fields W^1 and W^2 as a particle/anti-particle pair. The $W^\pm$ are therefore

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2) \tag{2.14}$$

The other physical bosons will be identified after the symmetry is broken. In §2.2.1, the Higgs mechanism was described for a $U(1)$ symmetry. In the electroweak sector, however, the symmetry to be broken is $SU(2)_L \otimes U(1)_Y$. The Higgs field Φ is therefore a weak isodoublet, with hypercharge +½. The T^3 = -½ component has zero charge and may therefore acquire a vacuum expectation value, v, as described in §2.2.1.

After spontaneous symmetry breaking (in the unitary gauge), the covariant derivative of Φ is

$$D_\mu\Phi = \frac{1}{\sqrt{2}}\left(\partial_\mu + \frac{ig}{2}\begin{bmatrix} W_\mu^0 & \sqrt{2}W_\mu^- \\ \sqrt{2}W_\mu^+ & W_\mu^0 \end{bmatrix} + \frac{ig}{2}\tan\theta_W B_\mu\right)\begin{bmatrix} 0 \\ v+H \end{bmatrix} \tag{2.15}$$

and the kinetic part of the Higgs doublet Lagrangian becomes

$$(D_\mu H)^2 = \frac{1}{2}(\partial_\mu H)^2 + \frac{g^2}{4}W^{+\mu}W_\mu^-(v+H)^2 + \frac{g^2}{8}(W_\mu^0 - \tan\theta_W B_\mu)^2(v+H)^2 \tag{2.16}$$

which may be interpreted as a mass term for the $W^\pm$ s, a mass term for a linear superposition of the real fields W^0 and B, and interactions between these massive fields and the scalar Higgs. The linear superposition of the W^0 and B fields is identified as the Z^0 boson

$$Z_\mu = \cos\theta_W W_\mu^0 - \sin\theta_W B_\mu \tag{2.17}$$

and (to 1st order) the $W^\pm$ and Z^0 masses are related by

$$\left(\frac{gv}{2}\right)^2 \equiv m_W^2 = m_Z^2\cos^2\theta_W \tag{2.18}$$

since there is a factor of half difference between the mass terms for complex and real fields. Corrections to this relationship from higher order perturbations result in an additional factor of 1.01. Finally, the orthogonal superposition to the Z^0 is identified as the photon, which has remained massless.

$$A_\mu = \cos\theta_W B_\mu + \sin\theta_W W_\mu^0 \tag{2.19}$$

Expanding the brackets in Eq. 2.16, and ignoring the kinetic and mass terms gives three and four-point interaction terms between the Higgs and the vector bosons.

$$L_{int} = \frac{g^2v}{2}W^\mu W_\mu H + \frac{g^2}{4}W^\mu W_\mu HH + \frac{g^2v}{4\cos^2\theta_W}Z^\mu Z_\mu H + \frac{g^2}{8\cos^2\theta_W}Z^\mu Z_\mu HH \tag{2.20}$$

where the charge superscripts have been omitted for clarity (the W's are assumed to be in charge conjugate pairs). If we write the coefficients in terms of the W and Z mass, this

becomes

$$L_{int} = \frac{2m_W^2}{v} W^\mu W_\mu H + \frac{m_W^2}{v^2} W^\mu W_\mu HH + \frac{m_Z^2}{v} Z^\mu Z_\mu H + \frac{m_Z^2}{2v^2} Z^\mu Z_\mu HH \quad (2.21)$$

The vector bosons therefore couple to the Higgs with strength proportional to their mass. Cubic and quartic self interactions of the Higgs, as shown in Figure 2.1, are also present, with strengths proportional to $3m_H^2/2m_W$ and $3m_H^2/4m_W^2$ respectively.

2.2.4 Fermion Masses

Although spontaneous symmetry breaking has been successfully used to explain the mass of the W and Z bosons, the fermions in this electroweak theory remain massless. An explicit mass term is not allowed because it would mix left and right handed states, and the theory is dependent on separate treatment of the two. However, interactions between the left-handed fermion doublet, the scalar doublet and the right-handed singlet are allowed. These are known as Yukawa interactions and have the form

$$L_{Yukawa} = g_l(\bar{\psi}_L \Phi \psi_R + \bar{\psi}_R \Phi \psi_L) \quad (2.22)$$

In the unitary gauge, this gives

$$\frac{g_e v}{\sqrt{2}} \bar{e}e + \frac{g_e}{\sqrt{2}} \bar{e}He \quad (2.23)$$

which is clearly a mass term for the electron, together with the interactions of the electron and the Higgs. It is noted that in the absence of the right-handed neutrino, no coupling between the Higgs and neutrino can arise, and therefore no mass term for the neutrino. The minor adjustments that are needed to this theory in the light of recent results will be explained briefly in §2.2.6.

If g_e is defined in terms of the electron mass, then the electron-Higgs interaction term may be written

$$\frac{g m_e}{2m_W} \bar{e}He \quad (2.24)$$

So the Higgs coupling to the electron is proportional to the electron mass. All fermion masses must arise in this way, since all fermions feel the electroweak force. Therefore, all fermions couple to the Higgs with strength proportional to their mass. It is clear then that the dominant decay mode of the Higgs will be to the heaviest particle/anti-particle pair that is kinematically available.

2.2.5 Higgs Searches

Sufficient precision measurements of the electroweak sector have been made that the Standard Model may be overconstrained. Using data from LEP 1 and 2, SLD, NuTeV and the Tevatron (m_t), a fit of the model to the data may be performed. This is used as a consistency check, by comparing values obtained from the fit with direct measurements. It can also be used to make an indirect measurement of the Higgs mass. The result of such a fit is given in Figure 2.2, as calculated by the LEP electroweak working group [3]. The fit is particularly sensitive to certain parameters, so the fit is plotted for two different calculations of $\Delta\alpha^{(5)}_{had}$, the photon vacuum polarisation due to light quarks. The solid line uses a more experimentally driven calculation, while the dashed line uses an alternative calculation which is more theory driven.

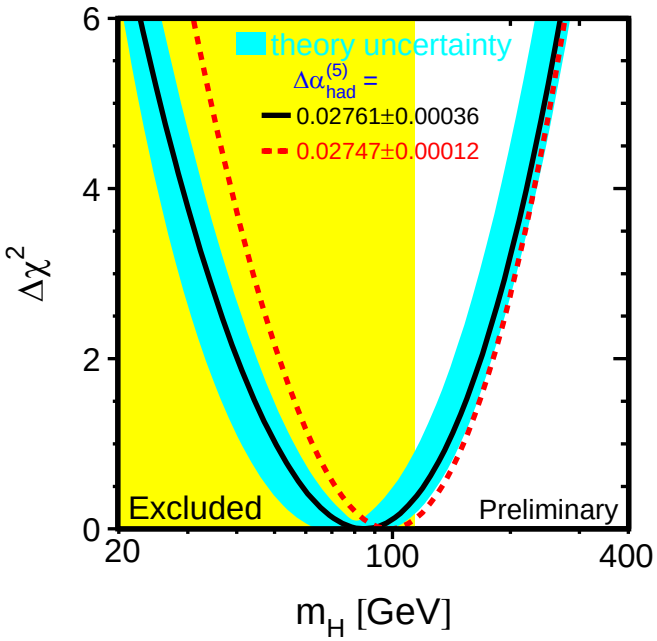

Figure 2.2 - Fit of precision electroweak data to log(m_H) (taken from [3]). See text for details.

The central value of m_H given by this calculation has already been excluded by direct searches at LEP. However, the 95% confidence level upper limit given by this calculation is 196 GeV (or 199 GeV if the theory driven calculation for $\Delta\alpha^{(5)}_{had}$ is used).

The absence of a signal in direct searches for the Standard Model Higgs boson at LEP 2 [4] indicates a lower bound on the Higgs mass, at the 95% confidence level, of 114.1 GeV. This region is indicated by the shaded box in Figure 2.2. The LEP searches look for a Higgs radiated by a Z^*, where the Higgs decays to $b\bar{b}$ or $\tau^+\tau^-$.

The search for the Higgs boson will continue with the Tevatron Run-II program. The main discovery modes here are via *WH* and *ZH* associated production, where the *W*/*Z* decays leptonically and the *H* decays to $b\bar{b}$, for $m_H < 140$ GeV, or to *W* pairs for $m_H > 140$ GeV. A 5σ discovery of the Higgs should be possible with 15 fb^{-1} integrated luminosity (obtainable by 2008) for $m_H < 125$ GeV. If the Higgs is not found at the Tevatron, the LHC will cover the remaining mass range. The favoured search channels for a light Higgs ($m_H < 150$ GeV) is in the di-photon decay mode. This channel will allow discovery at 5σ of such a Higgs with 30 fb^{-1} integrated luminosity, which should be obtained after one year of LHC running at low luminosity. For heavier Higgs bosons, decays to four charged leptons will provide the clearest signature. The LHC is expected to be capable of detecting Higgs bosons with mass up to a few hundred GeV, given several years of high luminosity running.

2.2.6 Neutrino Masses

As it stands, the Standard Model requires all flavours of neutrino to be massless. The recent measurements made at Super-Kamiokande [5] and the Sudbury Neutrino Observatory (SNO) [6] indicate that this is not the case. Both experiments show evidence for oscillations between neutrino flavours. This is only possible if the neutrino mass eigenstates are mixtures of more than one flavour, which has no meaning for strictly massless particles.

If the neutrino is treated as a Dirac particle like the electron, the introduction of a mass term mixes left and right-handed neutrinos. This is undesirable, because no right-handed neutrinos have been observed, nor left-handed anti-neutrinos. An alternative is to assume the neutrino is a Majorana particle. In this case, a mass term mixes neutrino and anti-neutrino states. This violates lepton number by ±2, and also implies that neutrinoless double beta decay is possible. If all fermion mass terms are assumed to arise from Yukawa interactions with the Higgs, Majorana fermions can be introduced by requiring the presence of a Higgs triplet, or two Higgs doublets arranged to transform as a triplet [7]. The LEP 1 measurement of the Z width has ruled out the most simple Higgs triplet model, but variants involving explicit lepton number violation or invisible Majoron models (the Majoron being the Goldstone boson associated with spontaneous breaking of lepton number symmetry) are still allowed.

Perhaps the most popular mechanism for explaining neutrino mass relies on the introduction of the right-handed neutrino as a singlet of colour, weak isospin and weak hypercharge. Such a particle is not involved in any interactions, and is therefore known as 'sterile'. A Majorana mass term for the right-handed neutrino may therefore be included in the Standard Model,

together with Dirac mass terms that mix the left and right-handed neutrinos. The Majorana mass of the left-handede neutrino is assumed to be zero. If the Majorana mass of the right-handed neutrino is sufficiently greater than the Dirac mass, then two mass eigenstates exist. The first is almost exclusively the right-handed neutrino and has large mass, and the second is almost exclusively the left-handed neutrino and has mass of order eV. This mechanism is known as the 'see-saw' mechanism, and allows the very small neutrino mass to arise in a more natural manner than the complex Higgs models associated with a Majorana mass.

Mass generation for neutrinos may have interesting implications for Higgs physics. In particular, as will be described later, invisible decays of the Higgs become more likely.

2.3 Supersymmetry

The Standard Model describes a great range of experimental data surprisingly well. However, it cannot be considered a fundamental theory. It has too many parameters to fully qualify as a theory, and it leaves a number of important questions unanswered. In particular, it does not explain why the electroweak breaking scale is so far removed from the only mass scale that can be constructed from fundamental constants, the Planck scale

$$M_P = \sqrt{\frac{hc}{G}} = 1.2 \times 10^{19}\ \text{GeV} \tag{2.25}$$

2.3.1 Gauge Hierarchy and Fine Tuning

A significant problem arises in the Standard Model when one considers radiative corrections to the particle masses. The one loop correction to fermion masses is found to be logarithmically divergent [8]

$$\Delta m_f \sim \frac{3\alpha}{4\pi} m_f \ln\left(\frac{\Lambda^2}{m_f^2}\right) \tag{2.26}$$

where Λ is some UV cutoff. This divergence is acceptable, because even at $\Lambda = M_P$, the correction is small. However, the one loop correction to the mass of scalar particles, such as the Higgs boson, is quadratically divergent [8]

$$\Delta m_H \sim O\left(\frac{\alpha}{4\pi}\Lambda^2\right) \tag{2.27}$$

and for any scale much greater than the electroweak scale (*i.e.* $\Lambda >> M_W$) the correction is large. In order to remove this divergence, we can set the bare mass m_H such that it cancels out the one loop correction, leaving a small mass of the order of the electroweak breaking scale. When calculated at the Planck scale, however, this cancellation must be accurate to 32

significant digits. Since the two loop correction will be proportional to $\alpha^2\Lambda^2$ these must also be cancelled. Though it is technically possible to cancel all the divergences arising from radiative corrections, increasingly accurate fine-tuning of the bare mass is required. This is hardly a satisfactory situation, and the problem is known as the gauge hierarchy or fine-tuning problem.

It is possible to resolve this problem by postulating a new set of particles. For each Standard Model fermion, a new boson is introduced, and for each Standard Model boson, a new fermion is introduced. Since the contribution to the mass correction due to a fermion loop has the opposite sign to the contribution from a boson loop, the quadratic divergence disappears [8]

$$\Delta m_H \sim O\left[\frac{\alpha}{4\pi}(\Lambda^2 + m_B^2)\right] - O\left[\frac{\alpha}{4\pi}(\Lambda^2 + m_F^2)\right] \sim O\left[\frac{\alpha}{4\pi}(m_B^2 - m_F^2)\right] \qquad (2.28)$$

This postulated fermion-boson symmetry is known as Supersymmetry (SUSY) and the partners to the Standard Model particles as superpartners, or SUSY particles. If SUSY is an exact symmetry, the fermion and boson in each pair must both have the same mass. However, no SUSY particles have been observed. If they exist, then SUSY must be a broken symmetry, and the sparticles must have greater mass than their Standard Model partners. The current lower bound from non-observation of sparticles at LEP 2 is around 90 GeV, depending on the particular sparticle, at the 95% confidence level (see §2.3.5 for details and references). In order to avoid the divergences of the hierarchy problem the sparticle masses must be no greater than the TeV scale.

2.3.2 Grand Unified Theories

It is generally accepted that the Standard Model is a low energy approximation to some fundamental theory. Such a theory is expected to bring about unification of all forces, known as Grand Unification. The coupling constants of the Standard Model (α_1, α_2, and α_3) are only effective constants, since they run with the energy scale. They can be parameterised using renormalisation group equations (RGEs), which allow the value of the coupling to be obtained at a given energy. In a GUT, we expect the three constants to converge at some scale (the GUT scale) where they are unified, and a single coupling constant takes over. However, in the light of precision electroweak data taken at m_Z by LEP 1 [9], the three Standard Model coupling constants do not converge on one point, as shown in Figure 2.3. However, if the RGEs are calculated with the addition of SUSY, the coupling constants may still converge, as indicated in Figure 2.4.

So, it is expected that a GUT must contain SUSY in some form, and that if SUSY is broken softly, it may resolve the gauge hierarchy problem. The next section outlines how a low-energy model of SUSY may be constructed.

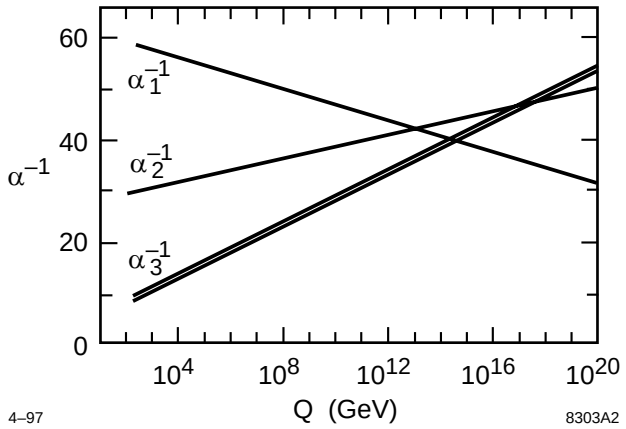

Figure 2.3 - The 3 Standard Model couplings as a function of energy scale Q, assuming only Standard Model physics (taken from [9]).

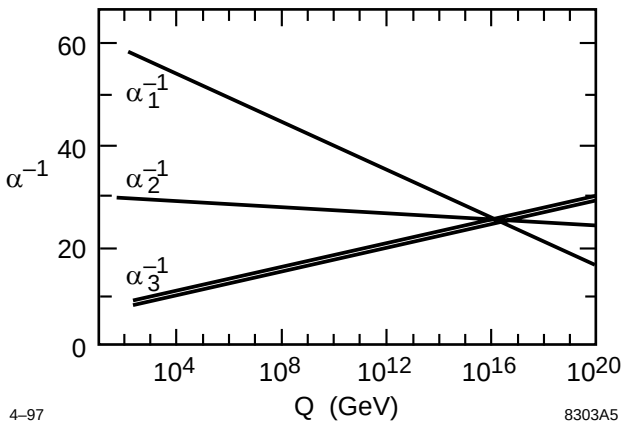

Figure 2.4 - The 3 Standard Model couplings as a function of energy scale Q, assuming SUSY in addition to Standard Model physics (taken from [9]).

2.3.3 The Minimal Supersymmetric Standard Model

When postulating a Supersymmetric theory, there are no restrictions on the number of extra particles one may add. It is usual, however, to limit the number of additional particles to the minimal number required to produce a Supersymmetric version of the Standard Model. The common features of such models are described in this section.

In the Minimal Supersymmetric Standard Model (MSSM) the Standard Model spin-½

fermions are each paired with a scalar superpartner, while the Standard Model gauge bosons are each paired with a spin-½ superpartner. These superpartners are known as the sfermions and the gauginos, respectively. Two Higgs doublets are required, one to generate mass for 'up type' quarks and one to generate mass for 'down type' quarks and the charged leptons. Each doublet contains a Higgs and a higgsino. As will be described later, the two Higgs doublets result in five physical Higgs bosons. A further complication results after electroweak symmetry breaking from the mixing between the neutral and charged gauginos and higgsinos, giving four neutralinos and four charginos. The particles of the MSSM are listed in Table 2.2.

particle	symbol	SUSY partner	symbol	mixing
Quarks (spin-½)	$\begin{bmatrix} u \\ d \end{bmatrix}_L \; u_R \; d_R$ $\begin{bmatrix} c \\ s \end{bmatrix}_L \; c_R \; s_R$ $\begin{bmatrix} t \\ b \end{bmatrix}_L \; t_R \; b_R$	Squarks (spin-0)	$\begin{bmatrix} \tilde{u} \\ \tilde{d} \end{bmatrix}_L \; \tilde{u}_R \; \tilde{d}_R$ $\begin{bmatrix} \tilde{c} \\ \tilde{s} \end{bmatrix}_L \; \tilde{c}_R \; \tilde{s}_R$ $\begin{bmatrix} \tilde{t} \\ \tilde{b} \end{bmatrix}_L \; \tilde{t}_R \; \tilde{b}_R$	
Leptons (spin-½)	$\begin{bmatrix} e \\ \nu_e \end{bmatrix}_L \; e_R$ $\begin{bmatrix} \mu \\ \nu_\mu \end{bmatrix}_L \; \mu_R$ $\begin{bmatrix} \tau \\ \nu_\tau \end{bmatrix}_L \; \tau_R$	Sleptons (spin-0)	$\begin{bmatrix} \tilde{e} \\ \tilde{\nu}_e \end{bmatrix}_L \; \tilde{e}_R$ $\begin{bmatrix} \tilde{\mu} \\ \tilde{\nu}_\mu \end{bmatrix}_L \; \tilde{\mu}_R$ $\begin{bmatrix} \tilde{\tau} \\ \tilde{\nu}_\tau \end{bmatrix}_L \; \tilde{\tau}_R$	
Gauge Bosons (spin-1)	W^α B g	Gauginos (spin-½)	$\tilde{W}^\alpha$ $\tilde{B}$ $\tilde{g}$	$\chi^0_{1,2,3,4}$
Higgs (spin-0)	$h \; H \; A$ $H^\pm$	Higgsinos (spin-½)	$\tilde{H}^0_{1,2}$ $\tilde{H}^\pm$	$\chi^\pm_{1,2}$

Table 2.2 - MSSM particles & sparticles. The physical states that result from

sparticle mixing are listed in the ‘mixing’ column.

In addition to the above particle spectrum, the MSSM Lagrangian contains every renormalizable supersymmetric interaction and, in the absence of a model for SUSY breaking, every soft SUSY breaking term. Soft SUSY breaking terms are defined as those which give mass to the superpartners, but do not reintroduce the hierarchy problem. The result is a model with 124 free parameters, most of which are due to the SUSY breaking terms. However, large regions of the MSSM parameter space have phenomenological properties that are inconsistent with experimental evidence.

In particular, measurements of the proton lifetime place a heavy constraint on allowed interactions. So far, the Lagrangian may contain terms that violate lepton number symmetry, and terms that violate baryon number symmetry. The presence of either one is acceptable, but the inclusion of both types of interaction is disastrous. Decays such as $p \rightarrow e^{+} \pi^{0}$ become possible at the tree level, and the proton lifetime is calculated to be $\Gamma_p \sim 10^{8}\ \mathrm{GeV}^{-1}$ [10]. This should be compared with the current limit of $\Gamma_p > 10^{56}\ \mathrm{GeV}^{-1}$ [10].

This problem can be solved by introducing conservation of *R*-parity, which is defined as

$$R = (-1)^{3(B-L)+2S} \tag{2.29}$$

where *B* is baryon number, *L* is lepton number and *S* is spin. All Standard Model particles have $R = 1$, while all SUSY particles have $R = -1$. The invariance of *R*-parity results in a predicted proton lifetime compatible with measurement, but also predicts that the lightest supersymmetric partner (LSP) must be stable and will not interact with Standard Model matter. This has important consequences for cosmology, as the LSP becomes a viable candidate for dark matter. Conservation of *R*-parity also has implications for SUSY signatures in collider physics. Any SUSY particle produced in collisions will decay, usually via some chain, to a stable state involving an LSP. Since the LSP is undetectable, an *R*-parity conserving SUSY signal will have a signature in the detector involving missing energy.

The phenomenological MSSM, even if *R*-parity is conserved, presents little progress over the Standard Model in terms of a fundamental theory. The 21 Standard Model parameters have been replaced by 124 MSSM parameters. However, if various reasonably well motivated assumptions are made, in particular about the unification of scalar masses at the GUT scale, the free parameters may be reduced to

$m_{1/2}$ - the gaugino mass parameter

m_0 - the scalar mass parameter

$\tan\beta = \frac{v_d}{v_u}$ - the ratio of Higgs doublet VEVs

A_0 - the scalar trilinear interaction strength

$\mathrm{sgn}\mu$ - the sign of μ, the Higgs doublet mixing parameter

This model is known as the constrained MSSM (CMSSM), and it is using this framework that collider physics potential is evaluated. Given a point in this parameter space, the SUSY particle mass spectrum is completely determined. However, in order to fully understand the phenomenology, a model for breaking SUSY is needed.

2.3.4 SUSY Breaking

If SUSY exists, it must be broken at some scale between the current experimental limits and the GUT scale. It seems natural to propose some Higgs-like mechanism, but this again leads to the hierarchy problem. However, SUSY can be broken by some force that is negligible at the GUT scale but becomes strong at some lower energy. Since the known particles do not feel such a force, it is assumed that SUSY is broken in some hidden sector. The effects of SUSY breaking are then mediated to the MSSM particles by some mechanism. Constraints on flavour changing neutral currents dictate that this mechanism must involve flavour independent interactions. There are therefore two obvious candidates; the interactions of the known gauge fields and their superpartners, and gravity. These will be explained in more detail later in this section. First, though, the special role played by gravity in SUSY theories will be described.

If SUSY is assumed to be a global symmetry, then spontaneous symmetry breaking will lead to a massless Goldstone boson, the goldstino. If, however, SUSY is a gauge symmetry, a suitable gauge field is required, and this gauge field will become massive when SUSY is spontaneously broken. Now, when forming a supersymmetric theory of gravity the spin-2 graviton, the gauge field of local coordinate invariance, is paired with a spin-3/2 particle called the gravitino. It turns out that the gravitino is in fact the gauge field required to maintain local invariance of SUSY. Hence the gravitino acquires mass when SUSY is broken, with important implications for phenomenology.

SUGRA

In models of SUSY breaking where supergravity mediates the broken symmetry from the hidden to the visible sectors, the SUSY breaking scale may be calculated to be of the order 10^{11} to 10^{13} GeV [11]. The gravitino will therefore acquire mass on the electroweak scale. However, since its couplings must be of gravitational strength, it plays no role in collider phenomenology. The effective LSP is likely to be the lightest neutralino or sneutrino.

GMSB

In gauge mediated SUSY breaking models, the interactions of the Standard Model gauge fields transmit broken SUSY from the hidden to visible sectors. These models contain a set of messenger particles that interact with both the hidden sector and the known gauge fields. When SUSY is broken in the hidden sector, the messengers acquire SUSY breaking terms, which are then transmitted to the visible sector via virtual exchange of the messengers.

In GMSB models of supergravity, the mediating effect of gravity will dominate if unsuppressed. In the case where gauge mediation dominates, the SUSY breaking scale is of the order 100-1000 TeV. The gravitino therefore acquires mass in the eV-keV range, and is the LSP.

Other models of SUSY breaking have been proposed. Some, such as anomaly mediated SUSY breaking, propose different methods of mediation between the hidden and visible sectors. Others suggest entirely new methods of breaking the symmetry, such as those involving extra dimensions on a variety of scales.

2.3.5 Experimental Limits on SUSY

It has already been described how the proton lifetime measurements place requirements on SUSY models, but a variety of other experimental results can also be used to constrain the parameter space. In general, the soft SUSY breaking terms of the MSSM produce large flavour changing neutral currents and large CP violation. In particular, the experimental upper limit on the flavour changing process $\mu \rightarrow e\gamma$ indicates that the different flavoured sleptons must have similar masses, to within one part in ~10^{-3} [11]. Also, measurements of the electric dipole moments of the neutron and electron place either upper bounds of ~10^{-2} on the size of the CP violating phases in the soft SUSY breaking terms, or require that some of the sparticles must have mass greater than ~1 TeV [11].

Direct searches for SUSY particles have been carried out at LEP and the Tevatron.

Interpretation of non-observation in terms of a mass limit is difficult, as it is highly model dependent. However, the preliminary results from the final year of LEP 2 exclude sleptons below 85 GeV [12], sbottom and stop below 94 GeV [13], and charginos below 92 GeV [14]. The neutral Higgs are excluded below 91.0 GeV [15], while the charged Higgs are excluded below 78.6 GeV [16]. A lower bound on the mass of the LSP can then be calculated at roughly 45 GeV [17]. All limits are at the 95% confidence level.

Measurements of the muon anomalous magnetic moment show a deviation from the Standard Model prediction of between 1.6 and 2.6σ [18]. Although this is insufficient evidence for new physics, the measurement is consistent with certain regions of MSSM parameter space. In particular, this result indicates preference for SUSY masses of a few hundred GeV and moderate tanβ.

Finally, various cosmological arguments can be used to place constraints on SUSY. These are based on arguments of the LSP as a dark matter candidate, baryogenesis etc. but are beyond the scope of this document.

2.4 An Invisible Higgs

In the Standard Model, the Higgs decays primarily to the heaviest particles that are kinematically available. Given the expected Higgs mass (115 - 200 GeV), this is likely to be a pair of *b* quarks or τ leptons, or possibly to a pair of *W* bosons. The branching fraction to invisible particles is very small, and can only proceed as

$$H \to ZZ^* \to \nu_l \bar{\nu}_l \nu_l \bar{\nu}_l \tag{2.30}$$

However, in extensions to the Standard Model, other invisible decays of the Higgs are possible. Low energy Supersymmetry is perhaps the most likely scenario in which such decays may occur with a significant branching fraction. Models involving large extra dimensions and the ‘see-saw’ mechanism for massive neutrinos may also give rise to an invisible Higgs, as could a fourth generation of heavy quarks and leptons. These mechanisms are briefly described below; a more detailed review of the motivation for studying an invisible Higgs is contained in [19].

2.4.1 Higgs Decays to SUSY LSPs

If *R*-parity conserving SUSY exists and the decay is not suppressed either kinematically or otherwise, the Higgs is expected to decay predominantly to a pair of LSPs. In SUSY models where the symmetry breaking is mediated by supergravity, the LSP is the lightest neutralino.

If it is kinematically available, any of the neutral Higgs may decay to a pair of these particles

$$h, H, A \rightarrow \chi_1^0 \chi_1^0 \tag{2.31}$$

For a moderate mass Higgs, in the 100-200 GeV range, these decays may form a significant branching fraction in certain areas of the SUSY parameter space. In particular, these decays are favoured by $\tan\beta$ of 5 to 10, and positive μ [20]. In other regions of SUSY parameter space the decay is either impossible or has a small branching fraction.

In GMSB models, the LSP is the gravitino, which has very small mass. However, the Higgs bosons only couple to the gravitino in conjunction with a neutralino or chargino. The lightest neutralino decays to a gravitino and a photon, so no decay chains involving purely invisible particles are possible.

2.4.2 Higgs Decays to Majorons

A number of extensions to the Standard Model, used to account for neutrino mass, involve the spontaneous breaking of global lepton number symmetry [21]. These models generally involve a more complicated Higgs sector than the simple doublet of the Standard Model. As mentioned in §2.2.6, the massless Goldstone boson associated with the symmetry breaking is known as the Majoron. In the allowed models, the Majoron interacts weakly with quarks, leptons and gauge bosons, but may interact strongly with the Higgs. If this is the case, the predominant decay mode of the Higgs will be to a pair of Majorons. Since the Majoron is undetectable, such decays will be invisible.

2.4.3 Higgs Decays in Large Extra Dimensions

It has been suggested recently that extra dimensions could become available on the TeV scale [22]. When neutrino mass generation is included in such models, invisible decays of the Higgs become possible. If the sterile right-handed neutrino is allowed to propagate in the extra dimensions, it acquires Kaluza-Klein excitations, denoted $\nu_R^{(i)}$. Although the decay of a Higgs to any given final state $\nu_L \bar{\nu}_R^{(i)}$ will be very small, a very large number of KK states is possible below the Higgs mass. In particular, if the number of extra dimensions is 3 or more, invisible decays of the Higgs to KK neutrinos may form a significant fraction of the total Higgs width. Other possibilities in such models involve the scalar Kaluza-Klein excitation of the graviton, known as the graviscalar, which is undetectable. In various cases, decays of the Higgs to such particles may dominate [23].

2.4.4 Higgs Decays to a 4th Neutrino

The precision measurements made by LEP-I at the *Z* mass have ruled out further lepton-quark

generations with light or massless neutrinos. Clearly, extra generations of very heavy particles are allowed, but it has also been suggested [24] that the measurements do not exclude a fourth generation if it contains a neutrino with mass greater than ~50 GeV that does not mix significantly with the first three neutrinos. If such a fourth generation exists, the predominant decay mode of a Standard Model Higgs with mass in the 100-150 GeV range would be to a pair of these neutrinos [25].

2.4.5 Invisible Higgs searches

Clearly, a comprehensive search for the Higgs boson should include the possibility that it decays invisibly. The LEP experiments all performed searches for an invisible Higgs, using the $l^+l^- + E_t^{miss}$ channel from Higgstrahlung production. Using data from all four LEP experiments and assuming a branching ratio to invisible states of 100%, the lower limit on the Higgs mass was found to be 114.4 GeV at the 95% confidence level [26]. The Run-II program of the Tevatron can be expected to make a 3σ observation in the $l^+l^- + E_t^{miss}$ channel for m_H = 125 GeV, provided the Higgs production rate is the same as the Standard Model and that it always decays invisibly [19].

Chapter 3 - The CMS Detector

This chapter describes a detector for high-energy particle physics; the Compact Muon Solenoid. The work described in this thesis was carried out as part of the experimental collaboration that is designing, building and will operate the detector. Where later chapters of this thesis require it, some aspects of the detector have been described in greater detail than others. In general, though, the entire detector is covered at the most basic level.

3.1 The Large Hadron Collider

The Large Hadron Collider is a new collider currently under construction at CERN. It will collide proton beams with a centre-of-mass energy of 14 TeV. The design luminosity of the machine is 10^{34} $cm^{-2}s^{-1}$, although it will initially operate at 2×10^{33} $cm^{-2}s^{-1}$. The bunch crossing rate is 40 MHz.

3.2 Overview of the CMS Detector

Although the design of CMS has evolved since its publication, the CMS Technical Proposal [27] gives a good overview of the detector and the principles behind its design. CMS follows a layered principle, similar to most other modern detectors at beam colliders. Figure A.1 (see Appendix A) shows the detector with a section removed to show its layered structure. The innermost layer consists of silicon pixel and microstrip detectors, for vertex and track reconstruction and measurement. These are followed by an electromagnetic calorimeter (ECAL) made of lead tungstate crystals. Beyond this, the hadronic calorimeter (HCAL) is a sampling calorimeter of plastic scintillator tiles inserted between plates of copper absorber. These detectors are all contained within a 4 Tesla magnetic field, supplied by a superconducting solenoid. The iron flux return of the solenoid is instrumented with various tracking detectors for muon detection and measurement. Beyond the flux return, and surrounding the beampipe, further calorimeters give access to very forward jets.

3.3 Magnet

A high magnetic field strength is required for good momentum resolution in the central tracker and muon chambers. The 13 m long $\times$ 5.9 m inner diameter solenoid used in CMS generates a high uniformity 4 Tesla field up to a pseudorapidity of $|\eta| = 2.4$. Other systems would require multiple magnets to achieve this coverage, making detector design more complex. The field within the central tracker is extremely uniform and will not require mapping. Hall probes will be used in the iron yoke to map the field. This data will be used to maintain good momentum

resolution in muon reconstruction.

3.4 Central Tracking

The central tracking detectors are essential for event reconstruction in CMS. Further details of the CMS Tracker design and performance are given in the Technical Design Report (TDR) [28] and an addendum to it [29].

3.4.1 Overview

The CMS central tracking system comprises silicon pixel detectors in the centre, surrounded by silicon microstrip detectors. The combined system of pixels and strips is used to reconstruct particle tracks, whilst the pixel system also provides impact parameter measurements and vertex reconstruction. A cross-section through the central tracking detectors is shown in Figure 3.1.

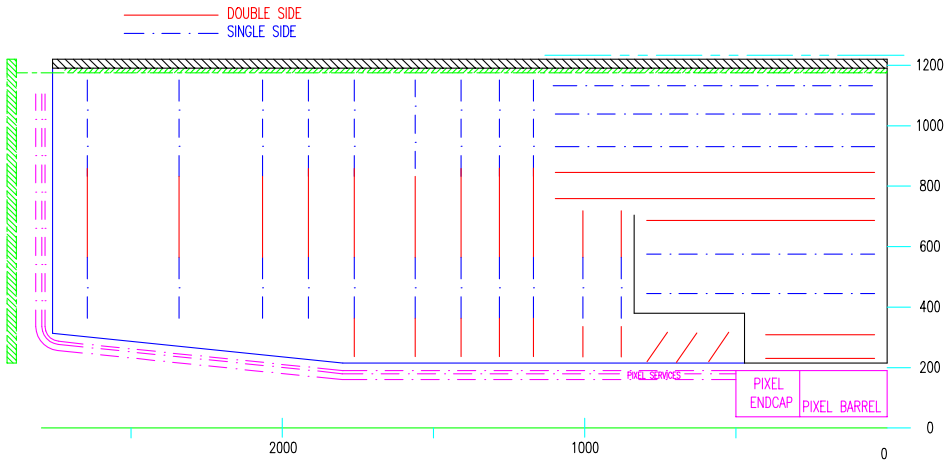

Figure 3.1 - Cross-section through the CMS Central Tracker

3.4.2 Silicon Pixel Detectors

The pixel detector assists in track reconstruction by providing 2 or 3 space points per track over the full rapidity range of the main tracker. It improves the impact parameter resolution for *b*-tagging and, by providing a much improved *z*-resolution in the barrel, allows 3-dimensional vertex reconstruction.

As shown in Figure 3.2, the pixel detector comprises two barrel layers, and two disks in each endcap. The initial configuration may also have an extra layer of detectors in the barrel at r = 4 cm. It is unlikely, however, that this layer will withstand high luminosity running, and may be removed before the LHC luminosity upgrade.

The basic detector unit used in the barrel has a sensitive area of 1.6 (r, ϕ) $\times$ 6.4 (z) cm, corresponding to a 128×512 array of 125 μm square pixels. These units are arranged in cylinders, with overlap in r and ϕ for hermeticity. In the endcap disks, the basic module is a trapezoidal carbon-fibre frame, with detectors on each side. The detectors overlap with those on the reverse in the r direction. The disks are made up of modules arranged in a 'turbine' geometry, allowing overlap in ϕ.

The effective (r, ϕ)-resolution of the pixel detector, assuming conservative alignment errors, will be about 20 μm. The effective z-resolution will be about 100 μm.

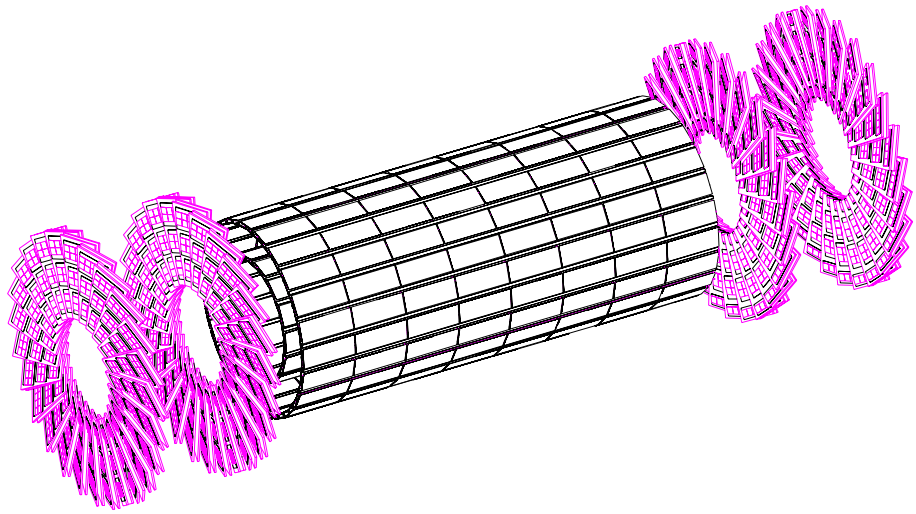

Figure 3.2 - The pixel detectors in high luminosity configuration

3.4.3 Silicon Microstrip Detectors

The radial region $20 < r < 120$ cm is instrumented with silicon microstrip detectors. Ten layers are used in the barrel region and eleven disks are used in each endcap, as shown in Figure 3.1.

The basic detector module used in the five single-sided barrel layers consists of two 300 μm silicon wafers, each with an active area of 62.5×62.5 mm^2. The two wafers are daisy-chained together, to form an effective strip length of 120 mm. The innermost layer will have a strip pitch of 60 μm, corresponding to 1024 strips. The pitch increases with each layer until layers 7 - 10, which have 446 strips on a 140 μm pitch. The five double-sided layers use modules consisting of four wafers, with the strips of the 2nd pair arranged at 100 mrad to the beam direction.

The basic module used in the endcaps consists of a wedge-shaped frame containing two wafers with radial strips. Each disk comprises up to four concentric rings. The outer (in *z*) disks have the inner (in *r*) rings removed, since the radiation environment is excessively harsh. The pseudorapidity range $|\eta| < 2.5$ is always fully instrumented. Rings 1 and 4 are instrumented with a second set of detectors, attached to the back of each module, allowing measurement of *r* in addition to ϕ.

3.4.4 Performance of the Central Tracker

The values for reconstruction efficiency and track resolution given below are the result of detailed simulation of the CMS tracker. For details of the simulation, and further results, see [28] and [29]. Results for muon track reconstruction, using both the central tracker and the outer muon chambers, are given in § 3.8.

In the central region of $|\eta| < 1.6$, high p_t isolated tracks can be reconstructed with a momentum resolution of better than

$$\frac{\sigma}{p_t} \approx (15 \cdot p_t \oplus 0.5)\ \% \qquad (3.1)$$

where p_t is measured in TeV. This gradually degrades with $|\eta|$ to

$$\frac{\sigma}{p_t} \approx (60 \cdot p_t \oplus 0.5)\ \% \qquad (3.2)$$

This resolution is well suited to reconstruction of narrow states decaying into charged particles, and is sufficient to ensure reliable charge assignment up to the highest kinematically available momenta.

The reconstruction efficiency for high p_t isolated tracks within $|\eta| < 2.5$ is greater than 90%. Tracks in jets with p_t in excess of 10 GeV/c are reconstructed with efficiency approaching 95%. Even tracks with momenta as low as 1 GeV can be reconstructed with > 85% efficiency.

The impact parameter resolution in the plane perpendicular to the beam is better than 35 μm for particles with p_t above 10 GeV, over the whole rapidity range covered by the tracker. The longitudinal impact parameter resolution is significantly better than 75 μm over most of this range.

It is expected that *b*-tagging efficiency of 50% can be achieved for jets with p_t between 50 and 200 GeV in the central rapidity region, with a mistagging probability of 1 or 2%. In the forward region, the tagging efficiency is expected to be around 40%, for the same mistagging

probability. Tagging efficiency is defined here as the proportion of *b*-jets that are correctly identified as such, while mistagging probability is the proportion of non *b*-jets that are incorrectly identified as being *b*-jets.

3.5 The Electromagnetic Calorimeter

The Electromagnetic Calorimeter (ECAL) will measure the energy of electrons and photons with high precision. A fully active scintillating crystal calorimeter provides the best performance, since most of the energy of the incident particle will be deposited in the crystal volume. Sampling calorimeters cannot attain the small stochastic contribution to the energy resolution that fully active calorimeters provide, and rarely achieve $10\%/\sqrt{E}$. This section gives a summary of the design and performance of the CMS ECAL taken from the TDR [30].

3.5.1 Lead Tungstate Crystals

The CMS ECAL uses crystals of lead tungstate ($PbWO_4$), which has a short radiation length (X_0 = 8.9 mm) and a small Molière radius (R_M = 21.9 mm). The low values for X_0 and R_M are both desirable; the former to limit the cost of the detector, the latter to limit the effects of pile-up. The barrel crystals are 23 cm (~ 26 X_0) long, to limit longitudinal shower leakage, and have front faces 22×22 mm^2 to match R_M. The endcap crystals are slightly shorter, since a preshower detector (described in §3.5.3) provides 3 X_0 of lead absorber. They also have a slightly larger front face to reduce the number of channels, since the cabling problem using barrel sized crystals in the endcap proved intractable.

Avalanche photodiodes (APDs) are used to collect the scintillation light of the barrel crystals. In the endcaps, however, the radiation dose is substantially higher, leading to increased leakage current in APDs and unacceptably high electronics noise. Vacuum phototriodes (VPTs) are therefore used for light collection of the endcap crystals.

3.5.2 ECAL Geometry

The geometric acceptance of the ECAL extends to $|\eta| = 3.0$, with precision energy measurement being carried out to $|\eta| = 2.6$. This limit has been set after an assessment of radiation dose and pile-up and matches the acceptance of the central tracker. The ECAL barrel covers the range $|\eta| < 1.479$, while the endcap covers the rest of the pseudorapidity range. A 3-dimensional view of the ECAL is shown in Figure 3.3. The lateral granularity of the crystals is $\Delta\eta \times \Delta\phi = 0.0174 \times 0.0175$ for the barrel, rising through the endcaps to $\approx 0.05 \times 0.05$.

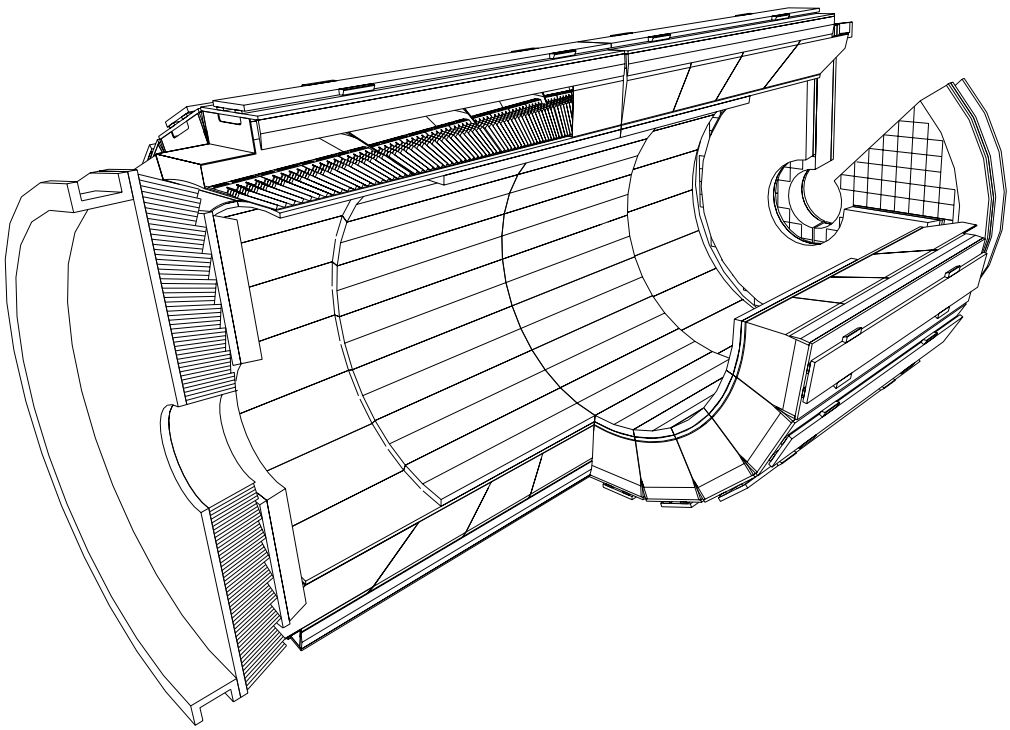

Figure 3.3 - 3-Dimensional representation of the ECAL

3.5.3 The Endcap Preshower

The function of the endcap preshower is to improve the distinction between single photons and neutral pions. In the barrel, photons from π^0 decay can be individually reconstructed and rejected. However, in the endcap, the π^0 is sufficiently boosted that the decay photons cannot be distinguished by the ECAL crystals alone, and appear as a single electromagnetic shower.

The preshower is placed in front of the ECAL endcap, covering the range $1.65 < |\eta| < 2.6$. The preshower uses two layers of lead converter ($2X_0$ and $1X_0$ thick), each followed by a plane of silicon strip detectors. The strips have 2 mm pitch, and the planes are aligned perpendicular to each other in η and ϕ. The impact position of the electromagnetic shower can be determined from the centre of gravity of the deposited energy, with a resolution of ≈300 μm at 50 GeV. A correction must be applied to the energy measured in the crystals to account for energy absorbed in the preshower.

3.5.4 Performance of the ECAL

For energies of 25 GeV to 500 GeV, the energy resolution of a single electromagnetic shower in the ECAL may be parameterized as

$$\frac{\sigma}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c \qquad (3.3)$$

where *a* is the stochastic term, *b* is the noise and *c* is the constant term. Expected values for the three terms in barrel and endcap are given in Table 3.1. These values are confirmed by test beam results and simulation [31]. Two values are given for the noise term, corresponding to low and high luminosity running (10^{33} and 10^{34} $cm^{-2}s^{-1}$), respectively.

	Stochastic / %	Noise / MeV	Constant / %
Barrel	2.7	155/210	0.55
Endcap	5.7	770/915	0.55

Table 3.1 - Electromagnetic calorimeter resolution parameters

The performance benchmark for ECAL energy resolution is taken as the di-photon mass resolution in $H \rightarrow \gamma\gamma$. It has been shown that, for a 100 GeV Higgs, this is 0.69 GeV at low luminosity with an overall reconstruction efficiency of 75% [32].

3.6 The Hadronic Calorimeter

The Hadronic Calorimeter (HCAL) acts in conjunction with the ECAL to measure the energies of particle jets and provides hermetic coverage for the missing transverse energy measurement. The main HCAL covers the pseudorapidity range up to $|\eta| = 3$ with a barrel and endcaps. Forward calorimetry outside the solenoid extends the hermetic coverage to $|\eta| = 5$, and will be covered in §3.7. Further details on the hadronic calorimeters in CMS are available in [33].

3.6.1 HCAL Detector

The main HCAL is a sampling calorimeter that surrounds the ECAL. The active elements are plastic scintillator tiles, sandwiched between layers of copper absorber. A tail-catcher outside the solenoid in the $\eta = 0$ region ensures adequate sampling depth. The light collection system consists of wavelength shifting fibres embedded in the scintillator tiles, read out by hybrid photodetectors (HPDs). The innermost of the 18 sampling layers is situated before any copper absorber, and is read out separately from the other layers. Since the hadronic energy deposited in the ECAL is underestimated by the crystal calorimeter, this facilitates improvement in the resolution of the combined ECAL + HCAL system.

3.6.2 HCAL Geometry

Like the ECAL, the HCAL coverage extends to $|\eta| = 3$. The barrel part of the HCAL covers the region $|\eta| < 1.4$, whilst the endcap covers the remaining pseudorapidity range. The lateral granularity of $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ for $|\eta| < 1.74$ is chosen to match the ECAL and

provides good di-jet separation and mass resolution.

3.6.3 Performance of the HCAL

As with the ECAL, the energy resolution of the HCAL may be parameterised using Eq (3.3). The expected energy resolution of the HCAL barrel is

$$\frac{\sigma}{E} = \frac{0.65}{\sqrt{E}} \oplus 0.05 \tag{3.4}$$

For the endcap it is

$$\frac{\sigma}{E} = \frac{0.83}{\sqrt{E}} \oplus 0.05 \tag{3.5}$$

For jet resolution, the performance benchmark is taken to be reconstruction of the di-jet decay of W and Z bosons. It has been shown [34] that the intrinsic di-jet mass resolution of CMS is about 19% at m_Z and 14% at 1 TeV, including the effects of initial and final state gluon radiation.

For the missing E_t measurement, the performance benchmark is taken to be the ability to discover and characterise supersymmetry. Studies using a fast parameterised Monte Carlo program have shown that CMS could discover squarks and gluinos up to masses of 2 TeV using a single charged lepton + jets + E_t^{miss} signature [35].

3.7 Forward Calorimeters

The main calorimeters extend to a pseudorapidity of $|\eta| = 3$. However, greater rapidity coverage is desirable for improved missing E_t resolution, and the reconstruction of very forward jets. The CMS detector uses 2 forward calorimeters (together known as the HF) outside the solenoid to extend the rapidity coverage to $|\eta| = 5$. Further details on the HF are given in [33].

3.7.1 Radiation Environment

The radiation environment in the forward region is extremely harsh. The innermost part of the HF ($4.5 < |\eta| < 5$) will receive a total particle flux of 6.0×10^6 cm^{-2} s^{-1} and the absorbed dose will reach values close to 100 Mrad/year. Showers in the HF will lead to large neutron fluxes and activation of the absorber. The detection technique employed must therefore be insensitive to neutrons and low-energy particles produced by the decay of radioactive nuclei.

3.7.2 Quartz Fibre Calorimetry

The HF consists of quartz fibres, embedded in copper absorber, parallel to the beam direction.

Cherenkov light is generated by particles traversing the fibres at a speed greater than that of light in quartz. In general, the only particles with sufficient speed are electrons. Since the amount of light generated is dependent on the incident angle of the electron, only those entering the fibre at $\approx 45° \pm 10°$ are considered [36]. A typical transverse shower size for both electromagnetic and hadronic showers is characterised by the Molière radius of the absorber. In the case of hadrons, the shower appears very short. As required, the technique is insensitive to neutrons since they carry no charge, and activation products if they lie below the Cherenkov threshold.

3.7.3 HF Geometry

The HF calorimeters will be cylindrically symmetric about the beam pipe and placed approximately 11 m from the interaction point. To cover the desired range of pseudorapidity ($3 < |\eta| < 5$), the active region of the calorimeter extends from r = 25 cm to r = 1.4 m. To achieve good sampling depth, the HF extends for 1.65m (or 10 nuclear interaction lengths) in the z-direction.

The transverse granularity of the HF is $\Delta\eta \times \Delta\phi = 0.17 \times 0.17$ over the full range of $3 < |\eta| < 5$. Whilst the calorimeter is constructed on an (x, y) geometry, the fibres are bundled according to an (η, ϕ) geometry. Different lengths of fibre are used to compensate for the differing response of the calorimeter to electromagnetic and hadronic showers. Full length (165 cm) fibres constitute the electromagnetic part, while short (143 cm) fibres constitute the hadronic part. 30 cm long ‘tail catcher’ fibres are used at the rear of the HF. Three photomultiplier tubes are used to detect the light from the three different lengths of fibre in each tower.

3.7.4 Performance of the Forward Calorimeters

The light yield in this type of calorimeter is extremely small; less than 1 photoelectron per GeV. The energy resolution is completely determined by fluctuations in this number. Beam tests [37] have shown the energy resolution for electrons to be

$$\frac{\sigma}{E_e} = \frac{1.5}{\sqrt{E}} \oplus 0.06 \tag{3.6}$$

whilst the result for pions is

$$\frac{\sigma}{E_\pi} = \frac{2.7}{\sqrt{E}} \oplus 0.13 \tag{3.7}$$

After optimisation of the weights apportioned to the three fibre lengths, a jet resolution of

$$\frac{\sigma}{E_{jet}} = \frac{3.0}{\sqrt{E}} \oplus 0.06 \quad (3.8)$$

can be expected [38].

3.8 Muon Detectors

The flux return of the magnet yoke is instrumented with muon detectors. The yoke itself acts as an absorber to reduce background from other particles and provides a magnetic field for momentum measurement. The muon system is designed to provide identification of muons up to $|\eta| = 2.4$, with a momentum resolution of 8-15% at 10 GeV (20-40% at 1 TeV) using the muon system alone. Matching tracks in the muon system with those in the central tracker will improve the momentum resolution to 1-1.5% at 10 GeV (6-17% at 1 TeV). Charge assignment should be correct at 99% confidence up to the kinematic limit of 7 TeV. Further details on the design and performance of the muon system are given in [39].

3.8.1 The Muon Detectors

The muon system in CMS comprises 3 detector technologies: Drift tubes (DTs), Cathode Strip Chambers (CSCs) and Resistive Plate Chambers (RPCs).

A drift chamber is the natural choice for the relatively low magnetic field, low rate environment expected in the barrel. A tube as the basic unit provides natural protection against damage from a broken wire.

CSCs are chosen for use in the endcaps since they are capable of providing precise space and time data in the high magnetic field and high particle rate found in the endcap. These chambers use a gas gap containing anode wires and bounded by cathode planes, one of which is segmented into strips running perpendicular to the wires.

RPCs have a fast time response, comparable to scintillators. They can be sufficiently highly segmented to provide a good p_t measurement to the Level-1 trigger. They therefore constitute a fast dedicated trigger which can identify muon tracks and assign bunch crossing with high efficiency.

3.8.2 Muon System Geometry

The barrel system consists of four stations integrated into the return yoke of the magnet. The inner 3 stations comprise 60 chambers, whilst the outer station has 70. The basic sensitive

element of each chamber is a drift tube of 400 ns drift time. The 12 planes of drift tubes used in each chamber are organised into 3 independent subunits, each with 4 planes of parallel tubes. The first two subunits in each chamber are used to measure the ϕ-coordinate, whilst the third subunit is used to measure the z-coordinate. Within each subunit, the four layers are staggered by half a cell, facilitating measurement of track coordinate and angle without any external time tag.

The endcap muon system comprises four stations of trapezoidal CSCs, arranged in concentric rings around the beampipe. Each chamber contains 6 gas gaps, with wedge shaped cathode strips in the radial direction and anode wires in the tangential direction. The wires in the innermost chambers are arranged at a 10 degree angle to the tangent, to compensate for the higher magnetic field in this region. Each station has two or three layers of chambers, which overlap for hermeticity.

3.9 Data Acquisition

3.9.1 Overview of DAQ

The Trigger and Data Acquisition (DAQ) system must reduce the event rate from the LHC bunch crossing rate of 40 MHz to an acceptable rate for permanent storage (~ 100 Hz). Figure 3.4 shows a schematic of the Trigger and DAQ. The Level 1 trigger consists of dedicated hardware processors acting on a reduced granularity subset of data from the calorimetry and muon subdetectors. Data from the front end is stored in readout buffers until Level 1 has made an accept/reject decision. If Level 1 accepts the event, the builder network sends data from the individual detector readout buffers to a filter unit. The filter unit is a commercial processor, running fast Higher Level Trigger (HLT) software. The HLT implements a variety of selection algorithms, that make the final decision to permanently store the event. In this case, the full event data, along with data from Level 1 and HLT results, is sent to computing services for storage on tape.

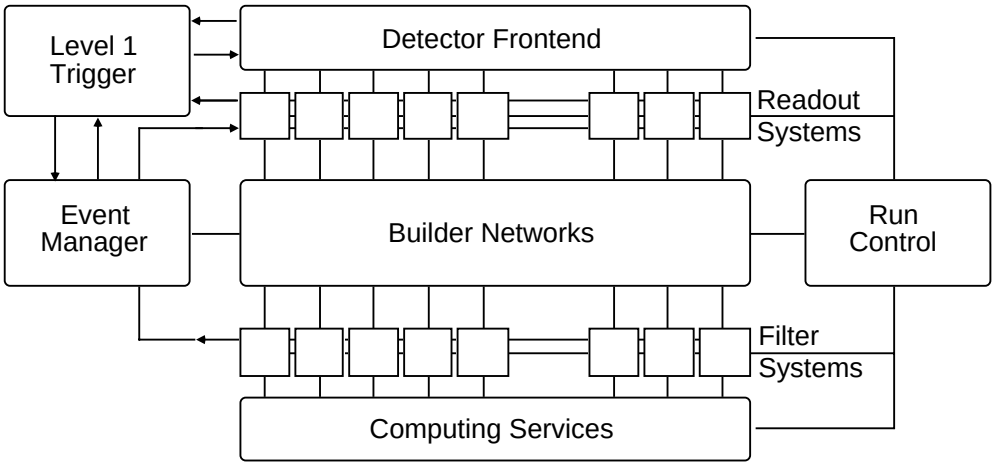

Figure 3.4 - Schematic of the CMS Trigger and Data Acquisition system

3.9.2 Higher Level Triggers

Once an event has passed the Level-1 trigger, a more complex series of trigger algorithms is applied. These algorithms will require greater granularity than is used at Level-1. This is achieved by using a large switch to send data from individual readout buffers to a node in a large farm of commercial processors (each node is known as a filter unit). The architecture is such that the full event data need not be sent immediately, but can be stored in buffers before passing through the switch when it is needed. This allows the HLT to use algorithms that can achieve some rejection factor with only, say, the calorimeter or muon system data. Events that pass each stage are processed with successively more complex algorithms requiring more and more of the event data. This allows a higher Level-1 accept rate for a given switch bandwidth. A possible example is to use only Muon system and Calorimeter data at Level-2, adding Pixel data at Level-3, before bringing in more Tracker data at Level-4. It is entirely possible that the full event data is not required until the event is submitted to tape.

3.10 The Level-1 Trigger

Since subsequent chapters are concerned with the hardware and algorithms of the Level-1 Calorimeter Trigger, this section is mainly devoted to a description of the baseline design of this system. A brief outline of the Muon Trigger is also given, together with some details of the Global Trigger. Further details of the baseline CMS Level-1 Trigger are available in the TDR [40].

3.10.1 System Overview

The Level-1 Trigger receives data from the ECAL, HCAL and muon readout systems. The

data is generally of coarser granularity in position (and energy in the case of the calorimeters) than is available offline. Level-1 receives this data at an event rate of 40 MHz (the bunch crossing frequency of the LHC). It is fully pipelined, with a total latency of less than 3 μs, including particle time-of-flight and cable delays. The Level-1 trigger architecture is shown in Figure 3.5. Candidate electrons, jets, taus and muons are found, and global energy sums are calculated. These 'trigger objects' are passed to the Global Trigger (GT) which can apply cuts to energy and position according to the Level-1 Trigger Table. If the event passes these cuts, a 'Level-1 Accept' signal is sent to the control system, which instructs the DAQ to pass the event to the filter farm.

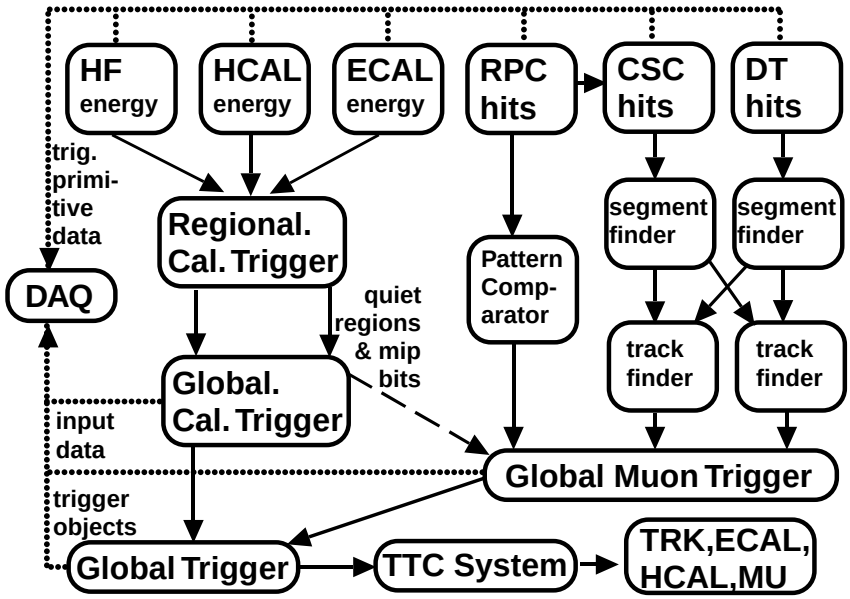

Figure 3.5 - Schematic of the Level 1 Trigger system

3.10.2 The Level-1 Muon Trigger

Muon candidates are found for each muon detector individually, in the DT, CSC and RPC triggers. This process starts by finding track segments from hits in individual CSC and DT detectors, then matching these to form tracks traversing the entire muon system. Some data sharing is necessary between the CSC and DT triggers to provide adequate cover of the overlap region. A separate algorithm is used for the RPC primitives, which forms track candidates directly. The candidates from each trigger system are then compared by the Global Muon Trigger using sophisticated algorithms incorporating p_t measurements, together with track quality criteria. The best four muon candidates are then passed to the GT for the final Level-1 decision.

3.10.3 The Level-1 Calorimeter Trigger

A schematic of the Level-1 Calorimeter Trigger is shown in Figure 3.6. For each bunch crossing, the trigger system requires digitized E_t values from all ECAL crystals and HCAL towers. The Trigger Primitive Generators (TPGs) convert data from the front-end electronics to transverse energy, and sum over constituent ECAL crystals or HCAL readout towers to obtain the trigger tower E_{t}. The TPGs also compute a 'fine-grain' bit for each ECAL tower, and a 'minimum ionizing particle' (MIP) bit for each HCAL tower (see §3.10.5 for details). Lastly, the TPGs assign these data to the correct bunch crossing. This allows the trigger to function synchronously in pipeline mode.

Next, 18 Regional Calorimeter Trigger (RCT) crates perform pattern recognition to identify electromagnetic showers in the ECAL. The RCT crates also sum ECAL and HCAL trigger towers as the first stage of the global energy flow algorithm. The Jet Cluster crate takes these energy sums, along with the HF trigger primitives, and performs pattern recognition to identify jets and hadronic τ decays. The electron/photon, jet and tau candidates are passed to the Global Calorimeter Trigger (GCT), along with the regional E_t sums. The RCT crates also generate 'quiet' and MIP data (again, see §3.10.5 for details), which is sent to the GCT.

The GCT sorts each stream of candidate trigger objects according to a rank value based on E_t. The top four candidates in each stream are passed to the Global Trigger for use in the final decision. The GCT also calculates the total scalar and vector transverse energy, and counts of jets above various thresholds. These data are also passed to the Global Trigger. The GCT collects the quiet/MIP bits from the RCT crates, and forward them, without processing, to the Global Muon Trigger.

Finally, the Global Trigger (GT) generates the Level-1 accept signal. This decision is based on highly programmable logic; transverse energy thresholds and position cuts can be applied to the incoming trigger objects, together with 'delta conditions' between trigger objects. Several such thresholds and conditions can be combined with boolean AND and OR logic to produce a trigger. The final trigger table comprises 128 such combinations. If any of these conditions are met, the GT passes a Level-1 accept signal to the Trigger Timing and Control (TTC) system, and the event is passed to DAQ.

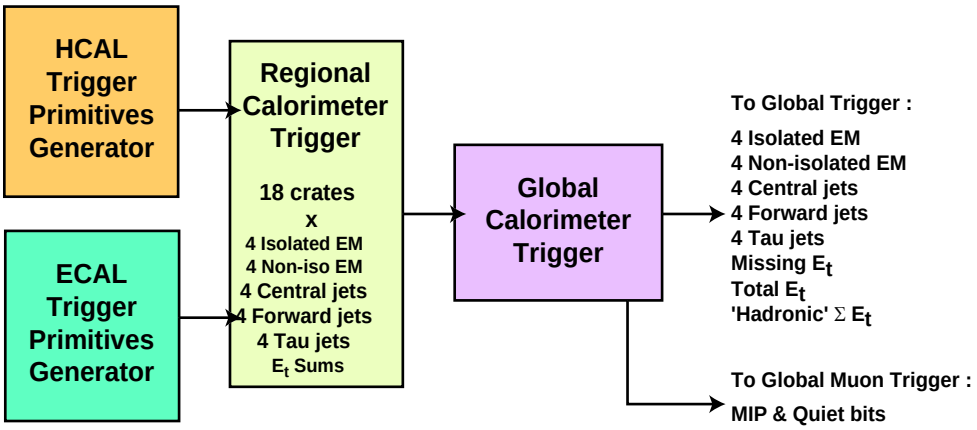

Figure 3.6 - Schematic of the Level-1 Calorimeter Trigger

3.10.4 Trigger Primitive Geometry

The main calorimeters are divided into 2448 barrel and 1584 endcap towers for trigger processing. In the barrel, the towers have lateral granularity $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$, corresponding to 5×5 crystals. In the endcaps, however the crystals are arranged in an x-y geometry, so the trigger towers do not follow exact (η, ϕ) boundaries. The average tower size is the same as the barrel up to $\eta \approx 2$. The number of crystals per tower varies from 25 at $\eta \approx$ 1.5 to 10 at $\eta \approx 2.8$.

The HCAL tower boundaries correspond to the trigger tower boundaries, except for $\eta > 1.74$, where the HCAL tower has twice the ϕ dimension of the trigger tower. Here, half the energy deposited in the HCAL is assigned to each of the two corresponding trigger towers.

High lateral granularity is less important in the forward calorimeters, since these detectors are not used in the electron/photon trigger. The baseline design groups $3\eta \times 2\phi$ readout towers into one trigger tower. The resulting $18\phi \times 4\eta$ segmentation is used in the jet and missing E_t triggers.

3.10.5 Trigger Primitives

The ECAL TPGs generate energy sums over the 25 crystals in each trigger tower. Each sum consists of a weighted sum of five consecutive time samples, to cater for the pulse-shape of the front-end pre-amplifier. The TPGs also perform bunch-to-bunch synchronisation, based on histograms of output data which can be compared with the LHC bunch structure.

The HCAL TPGs are based on the same principles as the ECAL TPGs. Since the pre-amplifiers used have a different pulse-shape to those used in the ECAL, the weights applied to

the time-samples are different.

Finally, the TPGs extract various fine structure data. These are represented as various bits associated with each tower. When summing over the 25 crystals in a tower, the ECAL TPG first sums the crystals in 5 strips along ϕ, corresponding to the direction in which the magnetic field bends charged particles. The ECAL 'fine-grain' bit is set when the highest energy adjacent strip pair has less than a programmable proportion of the tower energy. Approximately 98% of electrons and photons (both converted and non-converted) give a strip pair with more than 90% of the tower energy. For tower energies below the noise threshold, the fine-grain bit is set to zero.

The HCAL MIP bit is set if the energy in an HCAL tower is found in a window corresponding to a minimum ionizing muon. The bounds of this window are approximately 1.5 and 2.5 GeV.

3.10.6 Electron Photon Algorithm

The electron/photon algorithm is based on a sliding window of 3×3 towers, as shown in Figure 3.7. The candidate E_t is taken as the sum of E_t in the central hit tower plus the maximum E_t of the 4 broad side neighbours. This provides a sharper efficiency turn-on than simply taking the E_t of the central tower.

The isolated candidates are defined by the passing of two shower profile vetoes. The first demands that the fine-grain bit of the hit tower is set. The second demands that the ratio E_t^{HCAL}/E_t^{ECAL} is less than some programmable threshold (e.g. 5%).

The non-isolated candidates require passing of two additional vetoes. First, all 8 nearest neighbour towers must pass the Fine-grain and H/E vetoes. The second requires that all towers in at least one five-tower corner must be below a programmable threshold (e.g. 1.5 GeV)

In each calorimeter region (4×4 trigger towers), the highest E_t isolated and non-isolated electron/photon candidates are separately found. The 14 candidates found in each stream in one trigger region (8×28 trigger towers) are then sorted by E_t. Four candidates are kept from each trigger region, resulting in 72 candidates from the entire calorimeter. These are sorted by E_t, and the highest four are used by the Global Trigger in the final Level-1 decision.

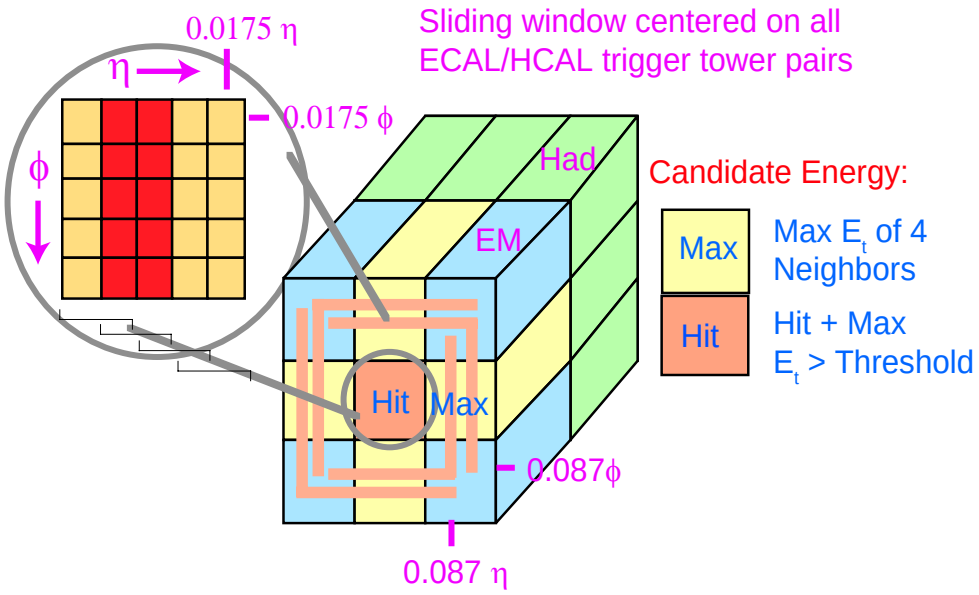

Figure 3.7 - The Level-1 Trigger Electron/photon algorithm

3.10.7 Jet Algorithm

The Level-1 jet algorithm is based on E_t sums over 4×4 trigger towers in the barrel and endcap. These 4×4 sums are referred to as jet towers. In the forward calorimeters, the trigger towers themselves are used.

The jet E_t is taken as the sum over a sliding 3×3 jet tower window. This provides seamless jet coverage up to $|\eta| = 5$. The central tower is required to have higher E_t than any of the neighbouring 8 jet towers.

In the central (barrel and endcap) region, an additional jet stream is provided, to enable lower thresholds to be used for hadronic τ decays. For each jet tower, a τ-veto bit is set if there are more than two active ECAL or HCAL towers in the 4×4 region. A jet candidate is classified as τ-like if none of the nine jet tower τ-bits is set.

The E_t of the jet candidate is re-calibrated according to E_t and η using a look-up table, to compensate for the non-linearity of the HCAL. The jets are then sorted in a tree similar to that used for the electron algorithm, and the highest four central, forward and central τ jets are passed to the Global Trigger for use in the final Level-1 decision. The normal (i.e. non-τ) jets are sorted in two separate streams; above and below an $|\eta|$ cutoff.

In addition, provision is made for jet counting. Counts of jets above programmable thresholds, in programmable η regions are provided to the Global Trigger to allow triggering

on events with a large number of low energy jets.

3.10.8 Global Energy Flow

The baseline algorithm for the Level-1 missing transverse energy calculation is given in Eq(3.9)

$$\left|E_t^{miss}\right| = \sqrt{E_x^2 + E_y^2} = \sqrt{\left(\sum E_t \cos\phi\right)^2 + \left(\sum E_t \sin\phi\right)^2} \qquad (3.9)$$

The summation is carried out over jet towers. An extension of this algorithm has been suggested [41]. Here, E_x and E_y are corrected according to the final 12 central, forward and tau jets, to allow for the non-linear response of the calorimeter to hadronic activity.

3.10.9 Level-1 Global Trigger

The Global Trigger receives candidate muons from the GMT, candidate electrons (isolated and non-isolated), jets (forward and central) and tau-jets, together with the Level-1 E_t^{miss} measurement, from the GCT. The GT compares the set of input objects to a set of trigger conditions. If these conditions are met, the GT passes a Level-1 Accept (L1A) signal to the TTC system. The TTC then distributes this the L1A to the readout buffers and DAQ.

The GT logic consist of three types of condition. The first consists of p_t or E_t thresholds and windows in η and ϕ that are applied to individual objects. The second consists of $|\eta|$ and $|\phi|$ differences between objects of the same type. Finally, $|\eta|$ and $|\phi|$ differences between objects of different type may be calculated.

These conditions are combined using AND-OR logic to form the final trigger table. This table has 128 entries, including all physics, calibration, testing, and others special triggers. Two examples of physics trigger conditions are given below.

For back-to-back electrons :

$E_t(1) > 20$ GeV AND

$E_t(2) > 15$ GeV AND

$0° < \phi(1) < 360°$ AND

$0° < \phi(2) < 360°$ AND

$170° < |\phi(1) - \phi(1)| < 190°$

For two electrons or muons together with missing transverse energy :

$$((E_t(\text{iso-electron 1}) > 15 \text{ GeV AND } E_t(\text{iso-electron 2}) > 15 \text{ GeV })$$
$$\text{OR}$$
$$(p_t(\text{muon 1}) > 15 \text{ GeV AND } p_t(\text{muon 2}) > 15 \text{ GeV }))$$
$$\text{AND}$$
$$E_t^{miss} > 100 \text{ GeV}$$

If any of the 128 trigger conditions are met, the GT initiates the Level-1 Accept process via the TTC system. Whilst the details of GT implementation are omitted here, and there are some restrictions on the set of possible trigger conditions, the information given here is sufficient for the purposes of this thesis.

3.11 Offline Computing and Software

Events stored by the DAQ consist of digitized detector data. If event rate stored by the DAQ is 100 Hz, and each event consists of ~1 Mbyte data, a single year of LHC running will produce over a Petabyte of data (1 Petabyte = 10^9 Mbytes). Clearly, analysis of this data will require very large computing facilities, running highly sophisticated software. The offline computing facilities, which will consist of large PC farms at several sites around the globe, will not be dealt with here. However, prototypes of the reconstruction and analysis software, together with Monte-Carlo data generation software, are relevant to the contents of this thesis, and are described below.

3.11.1 Software Overview

Figure 3.8 shows the software currently in use by CMS, together with the flow of data between packages. The software contains a mixture of Fortran (PYTHIA, GEANT, CMSIM) and object-oriented C++ software (Objectivity, COBRA, ORCA). The CMS-specific software is all at the prototype stage, and the final version software is likely to differ in many respects.

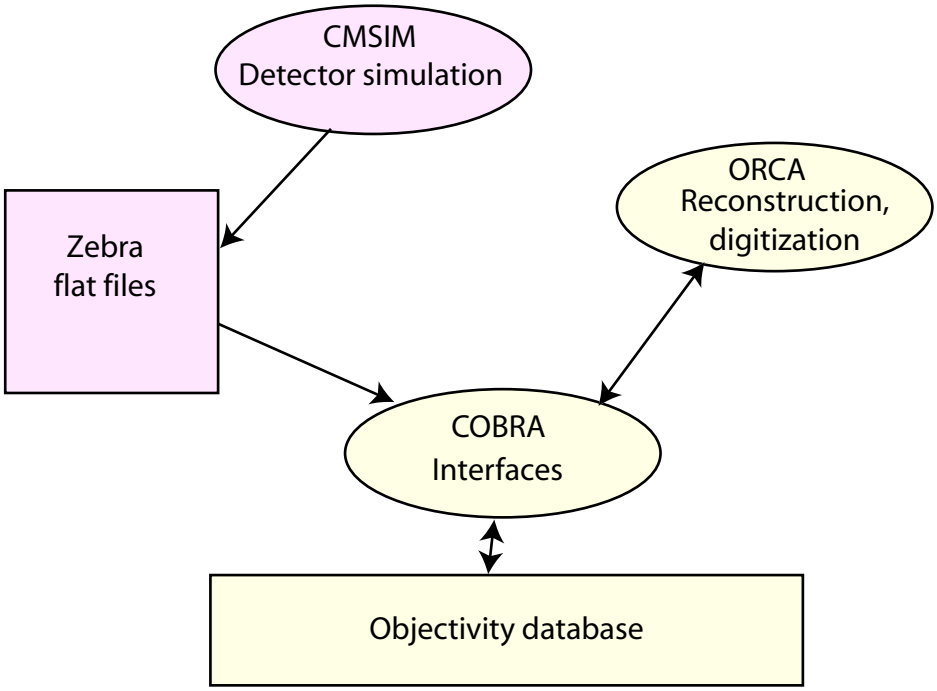

Figure 3.8 - CMS software

The interface between the reconstruction, or simulation, software and the database is provided by COBRA. When used with Monte-Carlo events, COBRA is responsible for simulating pile-up by constructing bunch-crossings from multiple events. It also provides access to the additional data stored with Monte-Carlo events (i.e. the paths and vertices of the generated particles).

The first stage of simulation is event generation. This consists of simulating a *pp* interaction, including the underlying hard-process, particle decay, fragmentation etc. The details of this process depend on the type of event (*pp* interaction) being simulated. Generally, PYTHIA 6.1 is used within CMS, however, any generator capable of producing HEPEVT ntuples can be used. In particular, SUSY events and special physics background signals use generators such as ISAJET and COMPHEP.

The passage of final state particles through the detector is simulated using CMSIM. This program is based on GEANT 3. This stage of processing includes the effect of the magnetic field, electromagnetic and hadronic showers, bremsstrahlung, etc. The output of CMSIM consists of the energy deposited in each element of the detector (known as the *hits*) for each input event. The hits are stored to file, which is later read by COBRA and stored in the database.

ORCA (Object-oriented Reconstruction for CMS Analysis) is used to reconstruct physics objects (electrons, muons, jets, etc.) from the detector data stored by the DAQ. It also contains code for simulation of the detector response (eg. photodetector response, electronic noise etc.) for use in production of Monte-Carlo data.

Chapter 4 - An FPGA Sort Processor

A fast sort processor is required for use in the CMS Global Calorimeter Trigger. The baseline solution is to use an Application Specific Integrated Circuit (ASIC) device. These, however, are extremely costly and offer little in the way of flexibility. A sort processor that can be implemented in a Field Programmable Gate Array (FPGA) was designed, based on a prototype ASIC. The design is described below, together with simulation results that indicate such a processor is a viable solution. On the basis of these simulation results, two technology demonstrators were used to test the design in hardware, also detailed below. The successful outcome of these tests has allowed the design of an FPGA-based GCT system to proceed.

4.1 The Global Calorimeter Trigger

The specification of the GCT given below is taken from the CMS Level-1 Trigger TDR [43], from which further details can be obtained.

4.1.1 Requirements

The Global Calorimeter Trigger (GCT) receives candidates trigger objects from the Regional Calorimeter Trigger (RCT) in five categories : isolated EM, non-isolated EM, central jets, forward jets and tau jets. These must be sorted, and the highest four candidates in each category passed on to the GT. Jets are counted before sorting, using several programmable E_t thresholds and $|\eta|$ regions. This input data is also used for luminosity monitoring.

In addition to the candidate trigger objects, the RCT also sends HCAL quiet bits, MIP bits and jet tower E_t sums (See §3.10.5 and §3.10.7 for definitions of these quantities). The E_t sums are used by the GCT to calculate the total scalar and vector transverse energy which are passed to the GT. The HCAL quiet bits and MIP bits are not processed, but merely forwarded to the Global Muon Trigger (GMT), where they are used for muon isolation.

Provision must be made in the GCT for capture of Level-1 data for storage to DAQ when a Level-1 accept decision is made. Finally, the system requires some interface with the outside world for control and testing purposes. The Trigger Timing and Control (TTC) system fulfills certain functions, supplying LHC clock and synchronisation data, as well as runtime control signals. An additional system is required for setup of the GCT, together with monitoring and error management.

4.1.2 GCT Technical Design

The original GCT design was based on ASIC processing chips [42]. However, it became clear in 1998/9 that the latest FPGA devices might well provide sufficient logic that they could be used to implement the data processors. This would offer a significant advantage over ASICs due to their flexibility and lower cost. FPGAs are highly flexible devices that can be programmed to perform a wide variety of tasks. This flexibility allows the GCT to use a single board design to fulfill all data processing requirements, reducing the development & production costs. If FPGA processors proved difficult to implement, the system design would be altered to accomodate the different processors required for sorting, summing, counting jets, and calculating $E_{\mathrm{t}}^{\mathrm{miss}}$ etc.

Data Type	Data Format	Multiplicity/BX	total bits/BX
Isolated EM	6-bit rank 4-bit position	18×4	720
Non-isolated EM	6-bit rank 4-bit position	18×4	720
Central jets	6-bit rank 5-bit position	9×4	396
Forward jets	6-bit rank 5-bit position	9×4	396
Tau jets	6-bit rank 5-bit position	9×4	396
Energy sums	11-bit E_t	36	396

Table 4.1 - GCT Input Data (excluding HCAL quiet/MIP bits)

The GCT input data is summarised in Table 4.1, with the exception of quiet/MIP data that is merely forwarded to the GMT. The data is received from the RCT on parallel ECL copper cables at 80 MHz. The connectors and input circuitry required to receive these signals place heavy demands on space. A dumb crate of Input Modules (IMs) is therefore used to receive and synchronise these signals, before transmission to a processor crate on high-speed serial LVDS cables. The HCAL quiet/MIP data is also received on the IM, but the high-speed serial LVDS output goes directly to the GMT.

Data Type	Data Format	Multiplicity/BX	Total bits/BX
Iso EM, Non-iso EM, Cen jet, Fwd jet, Tau jet	6-bit rank 9-bit position	5×4	300
Total E_t	13 bits	1	13
$\lvert E_t^{miss} \rvert$	13 bits	1	13
E_t^{miss} ϕ angle	6 bits	1	6
Jet counts	4 bits	8	32

Table 4.2 - GCT Output Data (excluding HCAL quiet/MIP bits)

The GCT output data is summarised in Table 4.2. As for input, the TPM transmits output data over high-speed serial LVDS links.

A schematic of the TPM is shown in Figure 4.1. Data inputs are on the left, each accompanied by a small FPGA (I/O) for synchronisation and a Dual Port Memory (DPM). During normal running, the DPM is used to store data for forwarding to the DAQ on receipt of a Level-1 Accept. It can also be used for injection and capture of test data, or monitoring purposes. The processing is carried out by four large FPGAs, arranged in a tree; input data goes to three ‘stage 1’ processor FPGAs and the output of these is sent to a single ‘stage 2’ processor FPGA. The final FPGA has several input/output ports attached. A single board CPU, accessed externally via Ethernet, is used for control.

One TPM has sufficient input data bandwidth to receive 72 trigger objects per bunch crossing. Five boards are therefore required to fulfill the sort requirements of the GCT assuming, of course, that one board has sufficient processing capability. Two further TPMs are used to perform the global energy flow calculations, count jets and perform the luminosity monitoring. A schematic of the boards required in the GCT processor crate is shown in Figure 4.2.

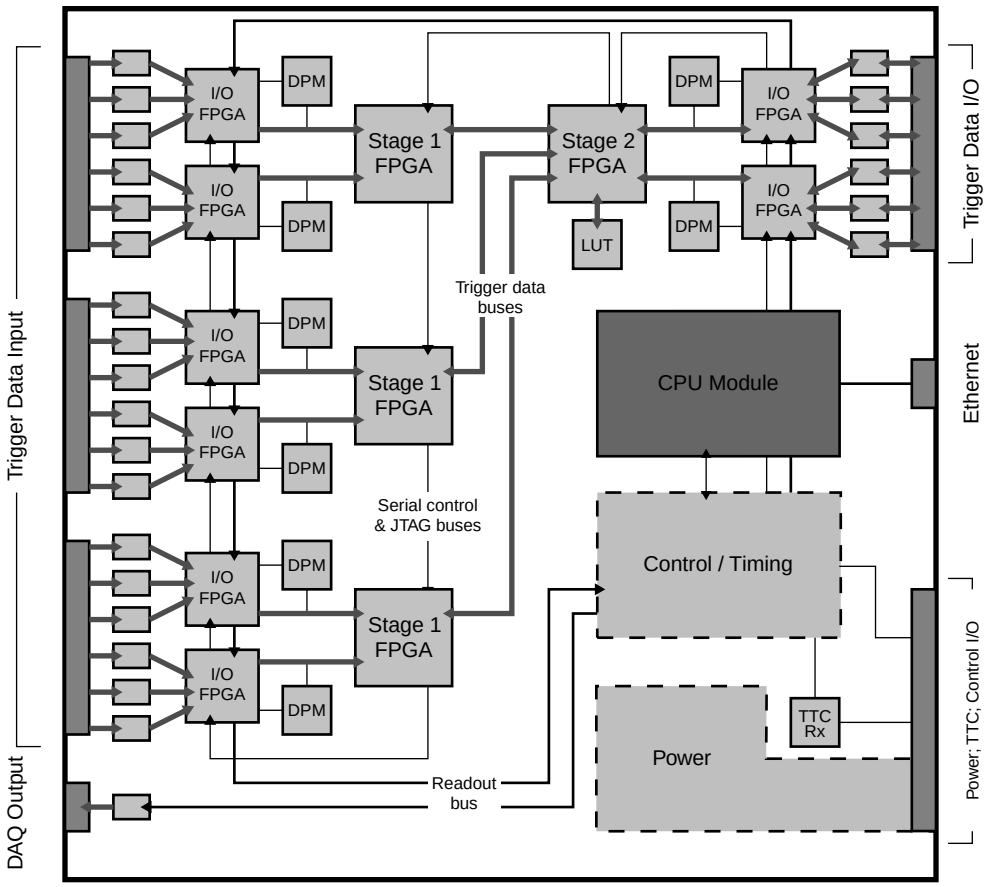

Figure 4.1 - Schematic of a GCT processor module (from CMS Trigger TDR [43])

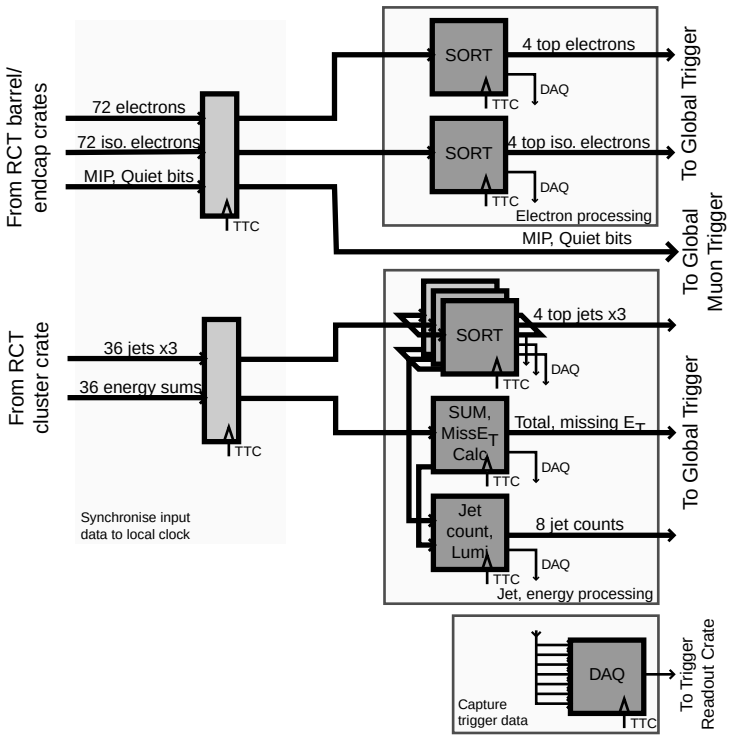

Figure 4.2 - Schematic of the GCT processor crate (from CMS Trigger TDR [43])

4.2 The Sort Algorithm

4.2.1 Requirements

At the input to the GCT, each trigger object is represented by a 6-bit rank word and a 4-bit position word. The rank value is a non-linear representation of the transverse energy of the object, and it is this value upon which sorting is based. The position word indicates the position of the object within the trigger region from which it was received. To fully identify the location of the object within the calorimeters, the GCT must add a further 5 bits to the position word.

Using the TPM design above, four sort processors must process the 72 input objects in each category of trigger object. The work is divided between the processors in a tree structure. The three 1st stage processors each receive and sort 24 input objects, sending the four highest rank (and therefore E_t) objects to the 2nd stage processor, which must then sort 12 objects and send the highest ranked four to the board output. It should be noted that the data width per object is greater at the output of the stage 1 processor than the input, and greater still at the output of the stage 2 processor, as position bits are added at each stage of sorting.

The sort processor detailed here is designed to fulfill the requirements of a stage 1 sort processor. It must receive 24 10-bit objects per bunch crossing, and output four 13-bit objects. The latency of the process should be as low as possible.

4.2.2 Algorithm

This algorithm described here is to U. Schäfer [44]. The implementation of this algorithm in an FPGA is described later, in 4.3.

The 1st stage sort nodes a required to find the highest four objects from 24 inputs. We will call these 4 output objects MAX_0, MAX_1, MAX_2 and MAX_3 in descending rank order. The input data is received as six groups of four objects. The objects in each group must be ordered. If this is not the case, additional processing is required upstream of the sort algorithm to achieve this. The first step of the algorithm is to sort the six leading objects, one from each group, and take the highest four. All comparators used in the algorithm are considered to have two inputs, X and Y, and a single 'X > Y' output. We therefore perform all 15 possible comparisons between the 6 leading objects, and may deduce their rank order from the results.

The group with the highest ranked leading object is labelled A. The objects within this group are labelled in descending rank order A_0, A_1, A_2 and A_3. The groups with the 2nd, 3rd and 4th placed leading objects are labelled B, C and D, respectively, and their contents labelled as for group A. We now know that

$$A_0 \geq B_0 \geq C_0 \geq D_0 \qquad (4.1)$$

and that

$$I_0 \geq I_1 \geq I_2 \geq I_3 \quad \text{where I = A, B, C, D} \qquad (4.2)$$

Table 4.3 shows the data after the first step of the algorithm. Any object in the table is greater than or equal to the object below it. Each object in the 1st row is also greater than or equal to the object to its right. It is clear that $MAX_0 = A_0$. We may also deduce that MAX_1, MAX_2 and MAX_3 must be contained in the subset shaded grey in Table 4.3 : {A_1, A_2, A_3, B_0, B_1, B_2, C_0, C_1, D_0}. All other objects can now be discarded, as they must be placed 5th or lower overall.

A_0	B_0	C_0	D_0	E_0	F_0
A_1	B_1	C_1	D_1	E_1	F_1
A_2	B_2	C_2	D_2	E_2	F_2
A_3	B_3	C_3	D_3	E_3	F_3

Table 4.3 - All input objects after the 1st step of the algorithm

It should be clear from Table 4.3 that MAX_1 must be either A_1 or B_0. Once this is determined, we may reduce the set of possible candidates for MAX_2. Similarly, once MAX_2 is known, we may select MAX_3 from another reduced set of candidates.

Table 4.4 shows the dependencies of MAX_2 and MAX_3 on the result of the selection of the higher outputs. In total, we perform 15 compares to determine all four output objects.The algorithm can be summarized in 4 steps :

- Compare 6 leading objects, one from each group
- Select the set $\{A_0, A_1, A_2, A_3, B_0, B_1, B_2, C_0, C_1, D_0\}$ according to the result of the previous step
- Perform 15 comparisons on the selected set
- Select MAX_0, MAX_1, MAX_2 and MAX_3

Note that if any two (or more) input objects have equal, non-zero rank values, then they appear at the output in an arbitrary, but consistent manner. If we label input groups from 0 to 5, then the compares made during the 1st step may arranged such that objects from a higher input group receive priority. This has been the case with all implementations of the algorithm.

.

MAX_0	MAX_1	MAX_2	MAX_3
A_0	A_1	A_2	A_3
			B_0
		B_0	A_2
			B_1
			C_0
	B_0	A_1	A_2
			B_1
			C_0
		B_1	A_1
			B_2
			C_0
		C_0	A_1
			B_1
			C_1
			D_0

Table 4.4- Output candidates. Candidates are selected in the order MAX_0 to MAX_3. (e.g. $MAX_0 = A_0$. If $A_1 > B_0$, $MAX_1 = A_1$. If $B_0 > A_2$, $MAX_2 = B_0$. If $C_0 > A_2$ and $C_0 > B_1$ then $MAX_3 = A_2$.)

4.3 FPGA Implementation

This section describes the physical implementation of the algorithm described above in FPGA technology for use in the GCT.

4.3.1 Algorithm Framework

Since the sort algorithm will be implemented in various types of chip during the prototyping process, it is desirable to separate the algorithm VHDL from surrounding logic, that may vary from chip to chip and board to board. Also, this 'glue logic' for a particular chip should be re-usable with other algorithms (e.g. the E_t^{miss} calculation).

The interface to the algorithm entity (henceforth referred to as the 'Algo') was kept as simple as possible. The 'glue' provides 40 and 80 MHz clocks, reset and sync signals to the algorithm along with input and output buses operating at 80 Mbit/s. Since the algorithms are

all required to process data continuously in pipeline mode, no control signals are neccessary.

4.3.2 Pre-sort

Each RCT output board sends four trigger objects, but does not order these objects. Since the algorithm described above requires ordered objects, an additional stage of processing is required in advance of the main sort. This ‘pre-sort’ is designed along similar lines to the main sort, but has only four inputs. Six such pre-sorts are required in each Stage 1 processor FPGA.

The two most obvious implementations of such a presort depend on input objects arriving in parallel, or serialised at 4 times the bunch crossing frequency (i.e. at 40 or 160 MHz). The first compares all input objects with each other (10 compares) and selects each object in parallel. The second consists of a chain of 4 stages. Each stage stores one object, and compares it with its input every cycle. If the input object is greater it is passed on to the next stage, otherwise the stored object is passed on, and its place is taken by the input object. Every fourth cycle, the stores are flushed to the output in parallel.

However, the speed of current FPGAs is insufficient to operate anything but the simplest designs at 160 MHz. So, this option is not considered further. A block diagram of the 40 MHz pre-sort is shown in Figure 4.3. The ‘4 of 4’ compare block consists of ten 6-bit comparators, together with logic to drive four one-hot encoded 4-bit select buses. These select buses drive four 10-bit, 4-to-1 multiplexers, that select the ordered objects. It is assumed that the compare and select may be performed in one clock cycle, so no additional pipelining is neccessary.

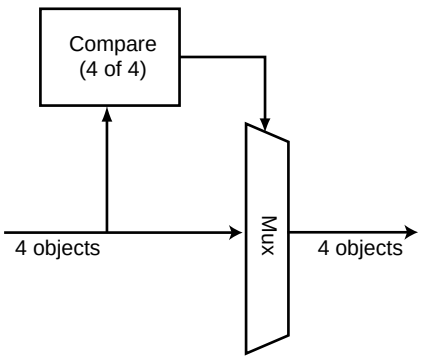

Figure 4.3 - Block diagram of the 40 MHz pre-sort.

4.3.3 Main Sort

As mentioned above, the stage 1 and 2 sort processors differ in number of inputs and size of

input object. The implementation described below is for a stage 1 processor (6 input groups, each of 4 objects). Extension to a stage 2 node is simple – unused inputs and associated circuitry are omitted, and extra bits are allotted for position information.

A block diagram of the main sort is shown in Figure 4.4. The '4 of 6' compare block finds the highest four of six 6-bit inputs. Fifteen 6-bit comparators perform a parallel compare of all possible two word combinations. The output of these comparators drive four one-hot encoded 6-bit select buses. This is essentially the same as the '4 of 4' block used in the pre-sort, with two additional input words. The one-hot outputs of the '4 of 6' block are used to drive ten 13-bit, 6-to-1 multiplexers, to select the objects shown in Table 4.3. Early simulation results suggested that the compare operation alone could only just be achieved in 25 ns, so a pipeline register is included between the compare and select blocks. Subsequent hardware tests and simulation have confirmed this. The select can be performed in less than 12.5 ns, so the next pipeline register is clocked on the falling edge of the 40 MHz clock.

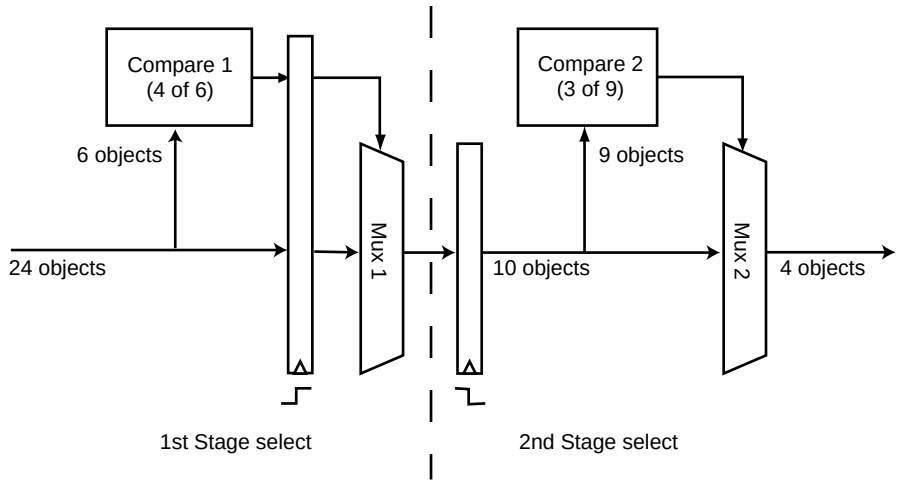

Figure 4.4 - Implementation of the sort algorithm.

The '3 of 9 find' module is similar to the '4 of 6', but some relationships between input words are already known (e.g., $A_1 > A_2$, $B_0 > C_0$ etc.). This reduces the number of comparators that are required, but the principle remains the same. The output words are encoded in the same way and are used to drive three 13-bit 9-to-1 multiplexers. These multiplexers select the MAX_1, MAX_2 and MAX_3 outputs; MAX_0 having been determined at the output of stage). This compare and select can be performed in less than 25 ns, giving the algorithm a total latency of 2.5 cycles of the 40 MHz clock.

4.4 Simulation

Before going to the expense of building test hardware, the feasibility of the sort processor design was checked with simulation software. The design was encoded in VHDL, and simulated at a functional level (i.e. with zero gate and routing delays) using Model Technology ModelSim. Exemplar Leonardo Spectrum was then used to synthesise the VHDL into an EDIF netlist and Xilinx Alliance software was used to 'place and route' the design, and calculate the gate and routing delays. ModelSim was then used to perform the gate-level simulation. The device files used for simulation corresponded to a Xilinx Virtex XCV300E-FG676-6.

4.4.1 Functional Simulation

Functional simulation was carried out for the presort and main sort blocks individually, and for a full stage 1 processor comprising 6 presort and a main sort. The simulation was carried out using a small set of test vectors, to check whether the VHDL conformed to expectations.

A timing diagram from the functional simulation of the presort using ModelSim is shown in Figure 4.5. The input signals are *data_in(0)* and *data_in(1)*. The four candidates per bunch crossing appear on these two lines over two cycles of *clk80*. The rank value of each candidate is shown first (as a 2-digit hex code), followed by the position value (a single digit hex code). The sorted candidates appear on signals *sorted_data(0)* to *sorted_data(3)* in descending order, and are multiplexed up to 80 MHz on signals *data_out(0)* and *data_out(1)*.

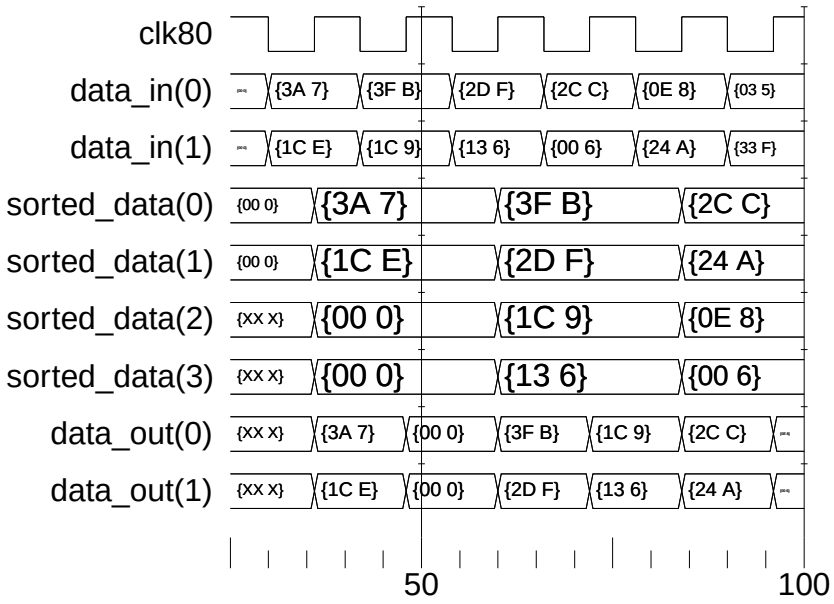

Figure 4.5 - Timing diagram obtained from functional simulation of the pre-sort. The horizontal scale is ns.

The timing diagram obtained from the functional simulation of the main sort algorithm is shown in Figure 4.6. As above, position and rank values are shown as two hex codes. The six input groups are labelled *Group 0* to *Group 5*, with the presorted objects appearing in descending order. The results of the 'Find 4-of-6' block are labelled *st1_sel*, and give the highest four of the six top rank objects, in ascending order. The four output objects appear at the bottom as *sort_out(0)* to *sort_out(3)*. The latency between the *st1_sel* signal and the output objects is one cycle of 40 MHz, so the output objects appear one column to the right of the corresponding input data. Note that the position codes of the output objects have an additional three bits, indicating which input group the object came from. It can be seen that *st1_sel* gives the correct values for each cycle, and that the output objects are indeed the highest four from each input set. The four highest input candidates from one set of input data are ringed with a solid line. In descending order, these are candidates A0, A1, B0 and C0. The D0 candidate is marked with a dashed ring, and the result of the 'Find 4-of-6' block is also ringed. As can be seen, the output data contains the correct candidates in the correct order, with the correct extra position code prepended.

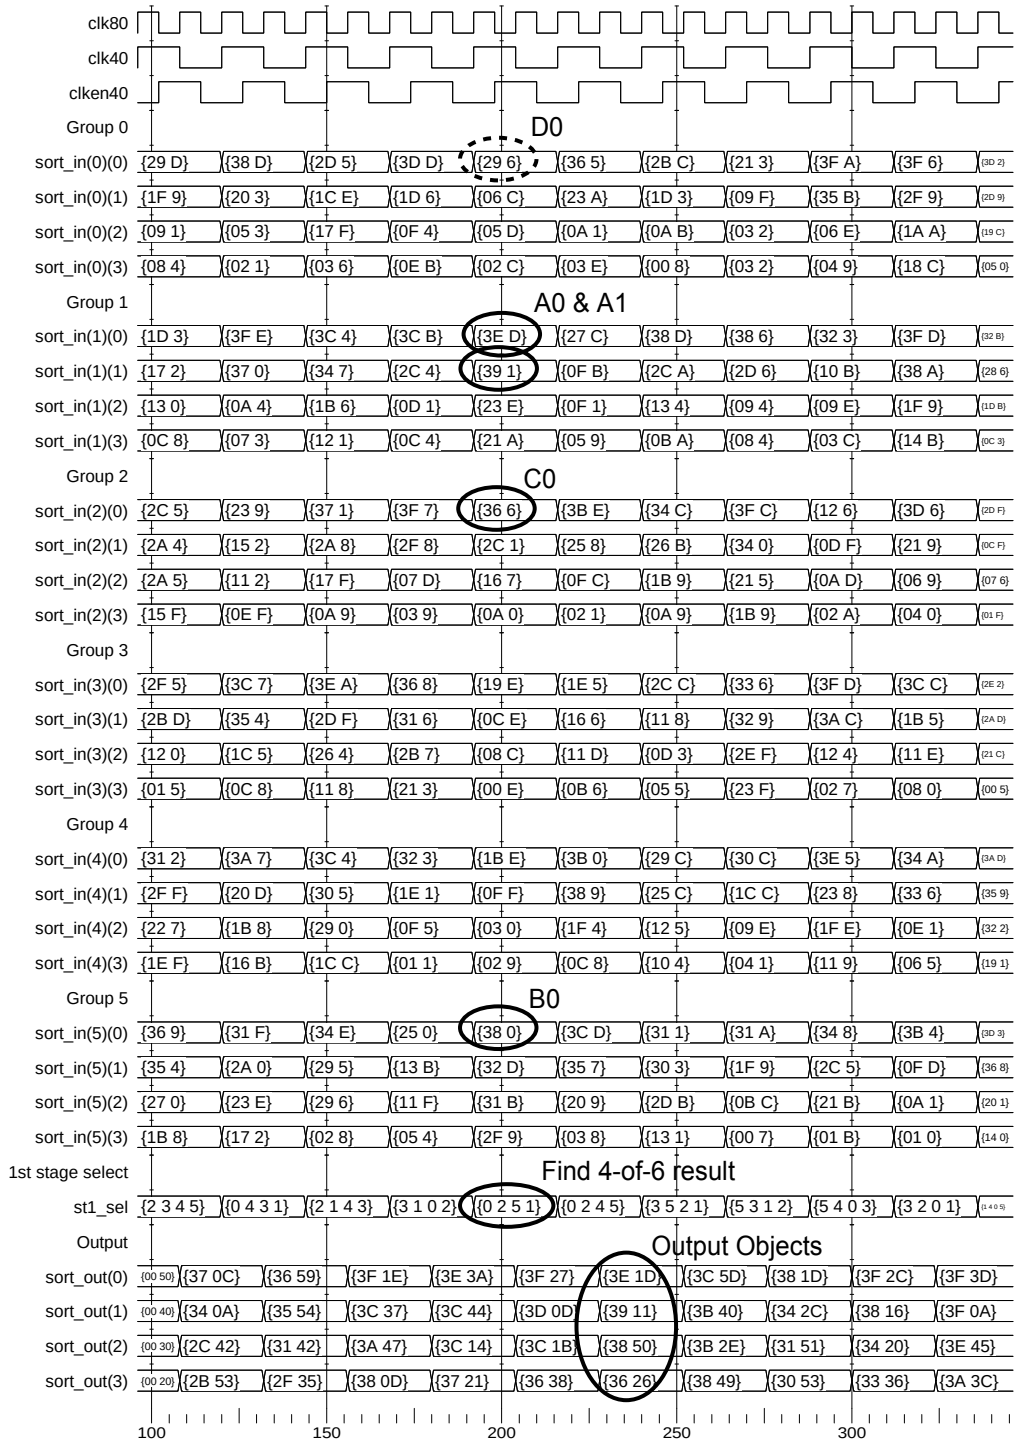

Figure 4.6 - Timing diagram obtained from functional simulation of the main sort. See text for explanation.

4.4.2 Gate-Level Simulation

A sample timing diagram from the gate-level simulation of the pre-sort is shown in Figure 4.7. Here, each object is represented by a 3-digit hex code. The first digit represents the position value of the input object, while the 2nd and 3rd digits represent the rank value. The input of the presort is driven by a register clocked on the rising edge of *clk40*. It can be seen that the pre-sort correctly sorts the four input candidates before the next rising edge of *clk40*, with ample safety margin. The output data is in fact valid after a maximum of ~13 ns. Clearly, if this figure could be reduced to <12.5 ns, the latency of the pre-sort could be reduced to half a cycle of the 40 MHz clock.

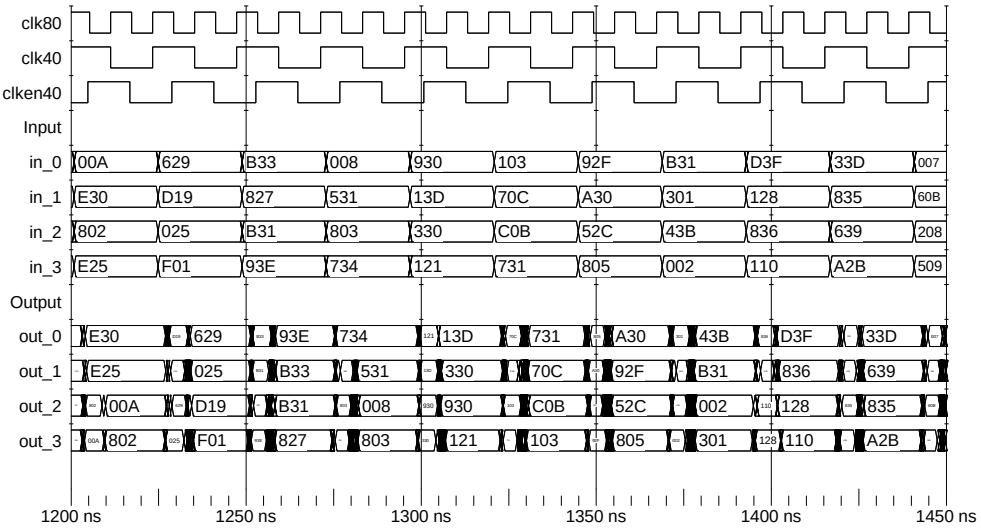

Figure 4.7 - Timing diagram for gate-level simulation of the pre-sort.

Figure 4.8 shows the timing diagram obtained with gate-level simulation of the main sort. These results were in fact obtained from a simulation of the full sort processor, including 6 presorts, but the signals shown are the input and output data of the main sort. As with the functional simulation, the six input groups are labelled *Group 0* to *Group 5*. Each input object is represented by a 3-digit hex code. The first digit is the position value; the 2nd and 3rd give the rank. An extra digit is added at the front of each output object, indicating which input group the it came from.

As can be seen, the input data are arranged within each group in rank order. The highest ranked object from each cycle of input data appears approximately 35 ns after the input data becomes valid, followed by the rest of the top four within 50 ns. The full output data is valid on the falling edge of *clk40*, 2.5 cycles after the input data is clocked in. The highest four input objects for one cycle are circled. It can be seen that these appear at the output below

(also circled).

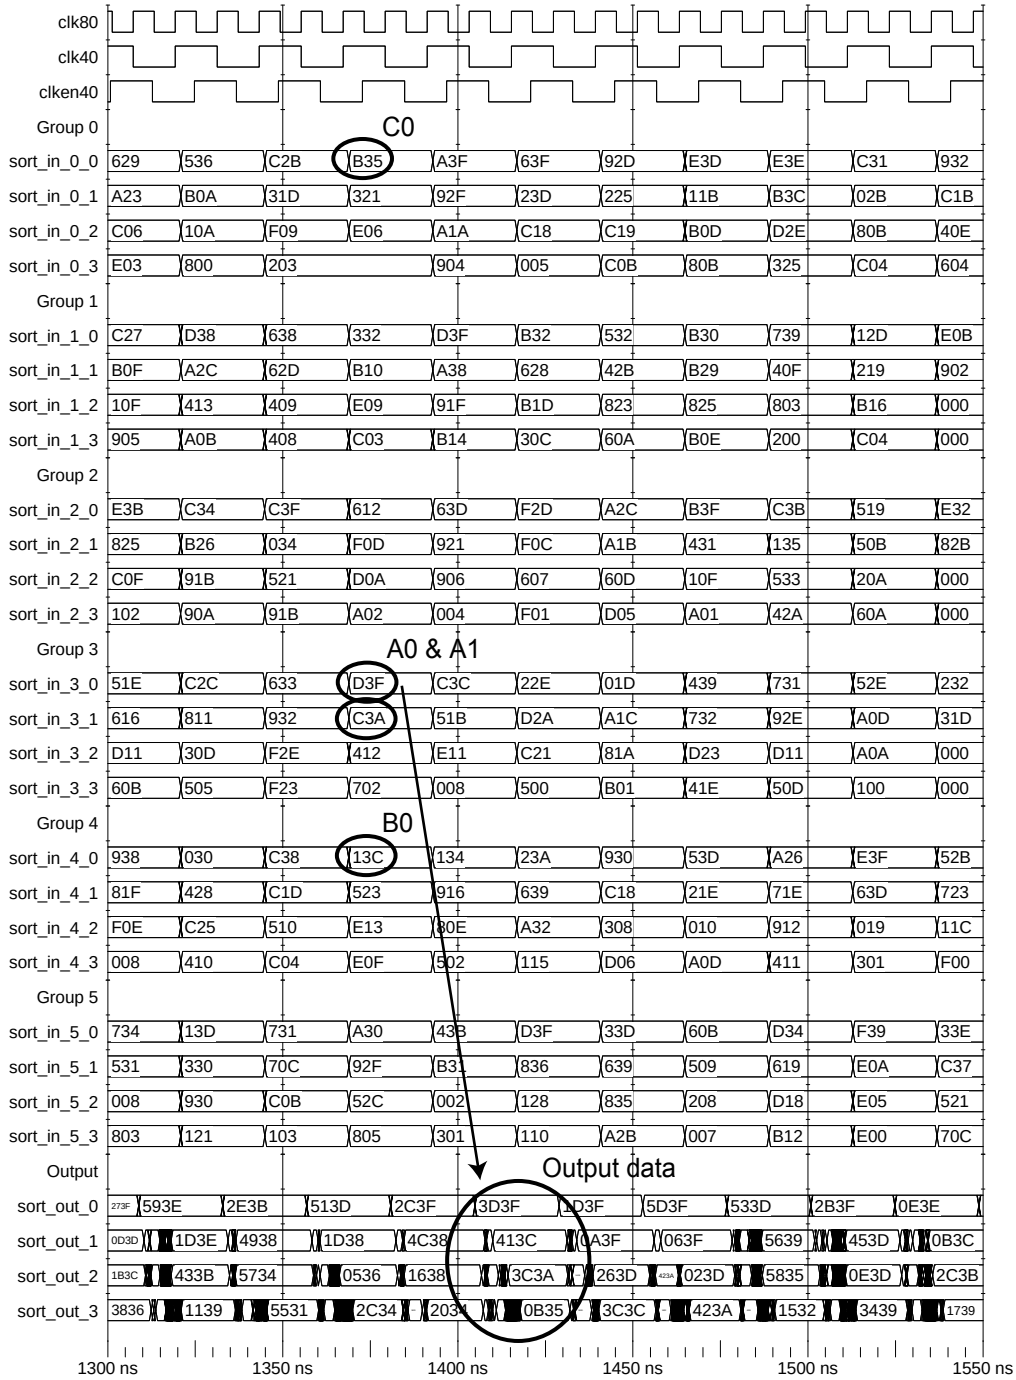

Figure 4.8 - Timing diagram for gate-level simulation of the main sort.

4.5 Virtex Test Platform

An FPGA test board was designed and built (by R.A.L. engineers) to test the sort processor. The test platform is shown in Figure 4.9, while Figure 4.10 shows a block diagram of the experimental setup. The board produces pseudo-random data using twelve 8-bit linear feedback shift registers (LFSRs). In this case, the D-input to each shift-register is the XOR of bits 5 and 6 (where the least significant bit is numbered 0). Each LFSR therefore produces a cycle of 255 8-bit values. The inital value is programmable via switches.

The LFSR data is fed into a Xilinx Virtex XCV300 FPGA. The output of this device is then sent via registers to an output port. The output data may be collected by a logic-analyser (in this case, a Hewlett-Packard HP16700A). The FPGA is programmed via a cable that interfaces the chip to a PC. The Xilinx software produces a ‘bitstream’ file from an EDIF netlist. After supplying power to the board, this file must be sent in serial mode to the FPGA. The FPGA is then ready to function as desired.

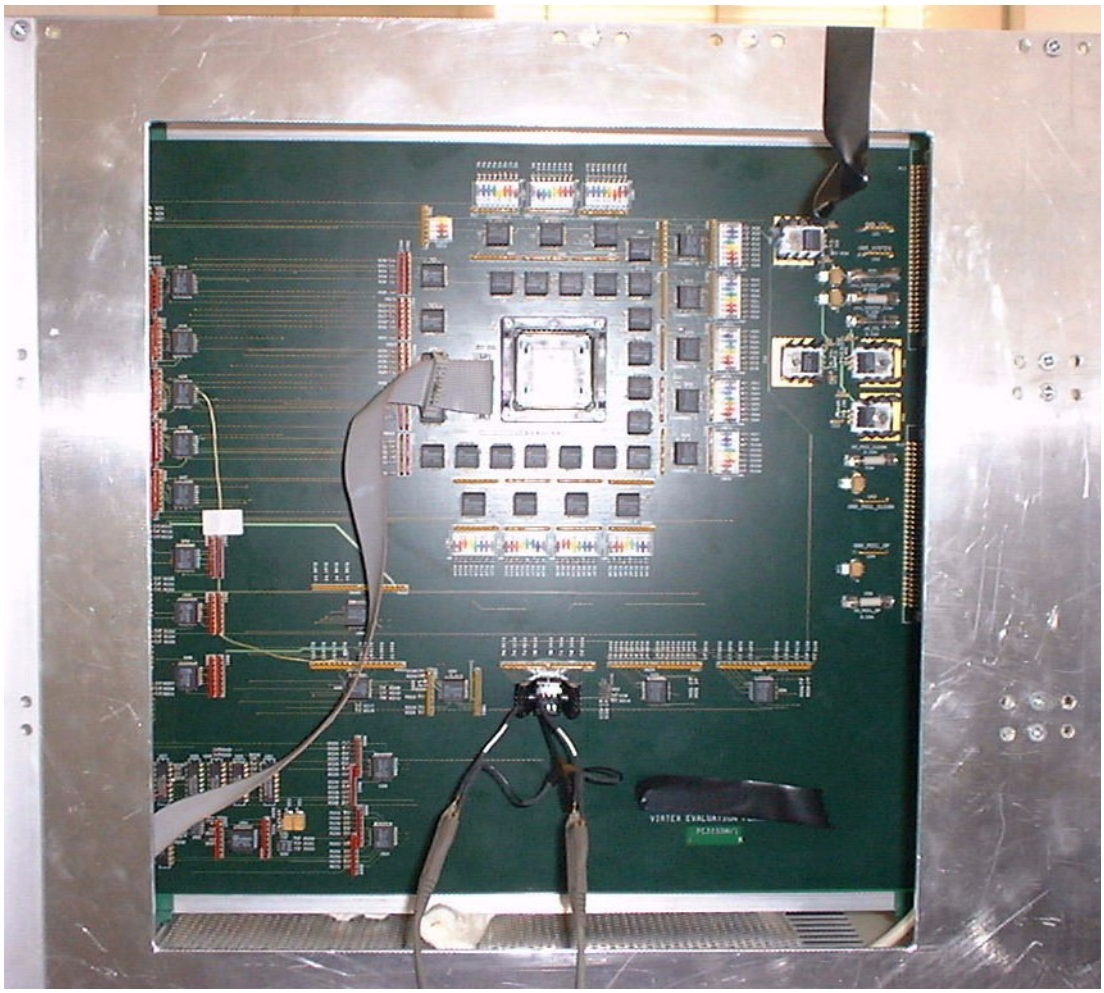

Figure 4.9 - The Virtex Test Platform

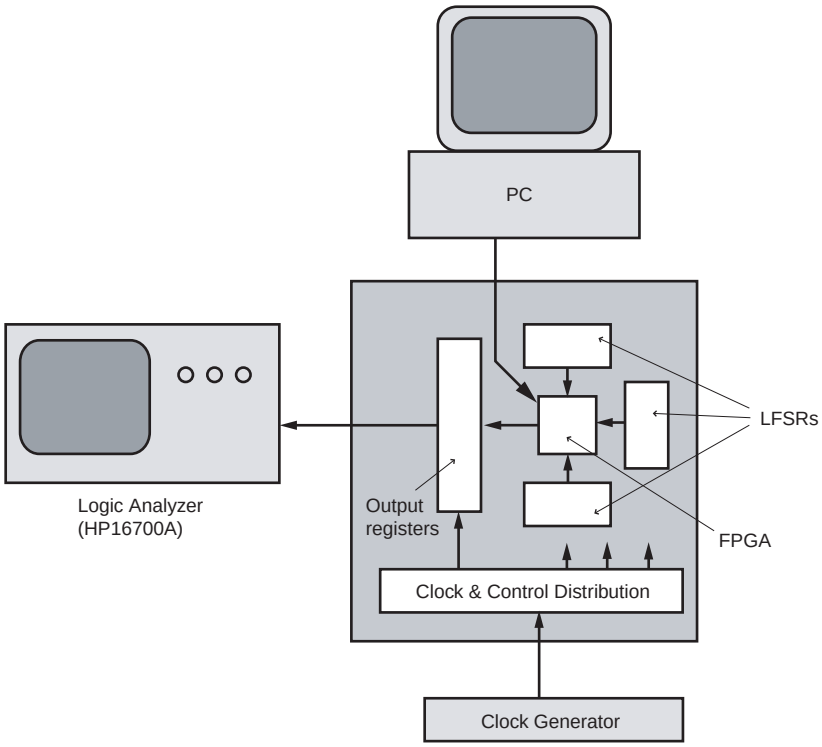

Figure 4.10 - External setup of the Virtex Test Platform

4.5.1 Board Tests

The board was thoroughly tested before use. The voltage regulators and clock distribution circuits were tested with an oscilloscope. Some dead clock paths were observed, but were repaired. The output data paths were tested by programming the FPGA with a simple counter circuit connected to the output pins. Observation of the expected patterns on the logic-analyzer confirmed the output circuits were functional. The LFSRs were then tested by programming the FPGA with ‘straight through’ connections from the LFSR inputs to the outputs. No dead output paths were observed. Several such programs were necessary, since the number of inputs outweighs the number of outputs. Again, observation of the expected LFSR patterns on the logic analyzer confirmed the LFSRs were functioning. Two dead input paths were observed, that could not be repaired.

4.5.2 Maximum data rate

The FPGA was programmed with binary counters and LFSRs, to investigate the maximum rate at which data could be captured. Only the FPGA, the output registers, and the logic analyser were used in this test. A sample of 512,000 cycles was captured at a clock speed and data rate of 160 MHz, and no discrepancies with prediction were observed. The first failures were observed around 180 MHz.

4.5.3 Proof of principle

The top-level ‘glue’ used for the Virtex test platform is relatively simple, as the FPGA has little more connections than the algorithm interface described above.

The FPGA was programmed with the sort design described in §4.3.2. Input to the sort came from the on-board LFSRs, and output was fed to the logic analyzer. The board design meant that synchronisation of the LFSRs was not possible, so spare lines were used to output one bit of each LFSR to the logic analyzer. The value of each input LFSR could then be calculated for any given clock cycle.

Critically, at this stage the HP16700 became temporarily unavailable, so a Thurlby-Thandar TA4000 was used in its place. This analyzer is only capable of synchronously capturing data below 50 MHz, so the data was read out at 40 MHz. At this rate, only two candidates per cycle could be read out. The FPGA was first configured to output the 1st and 2nd highest candidates. Comparison of output data with predicted results showed no discrepancies, suggesting the sort was working as desired. Next, the FPGA was configured to output the 3rd and 4th highest candidates. Again, comparison of output data with expected results showed no errors.

Whilst a test of the full design was not possible, the principle of sorting in a Virtex FPGA had been demonstrated.

4.6 Generic Test Module

A more sophisticated test board (the Generic Test Module, or GTM) was designed, with more of the functionality required of a GCT processor module, to demonstrate the full sort algorithm, including presorts, and the other GCT algorithms. A functional diagram of the GTM is shown in Figure 4.11, and a photograph in Figure 4.12. Processing can be carried out using two large Virtex-E FPGAs. Data paths are provided to Dual-Port Memories (DPMs), ‘daughter’ cards, and between the processor FPGAs themselves. A third FPGA provides interface to VME and board control signals. Clock signals are provided by a TTC prototype receiver.

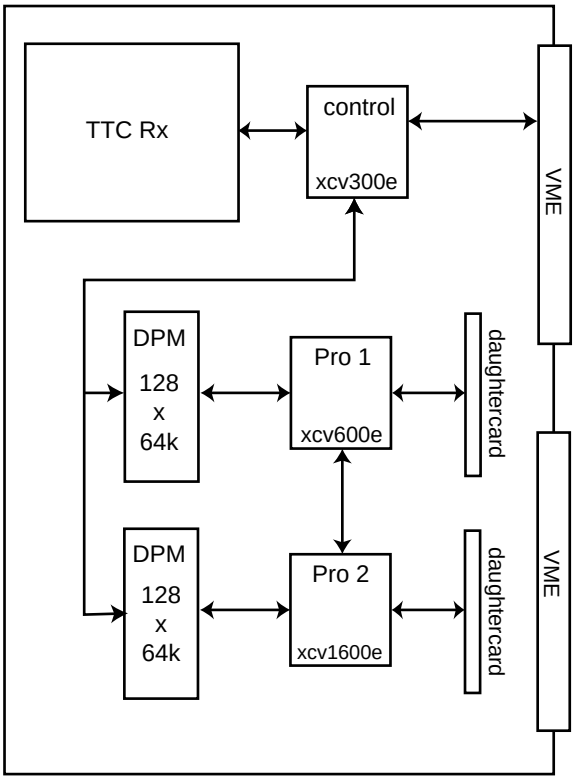

Figure 4.11 - Functional diagram of the GTM

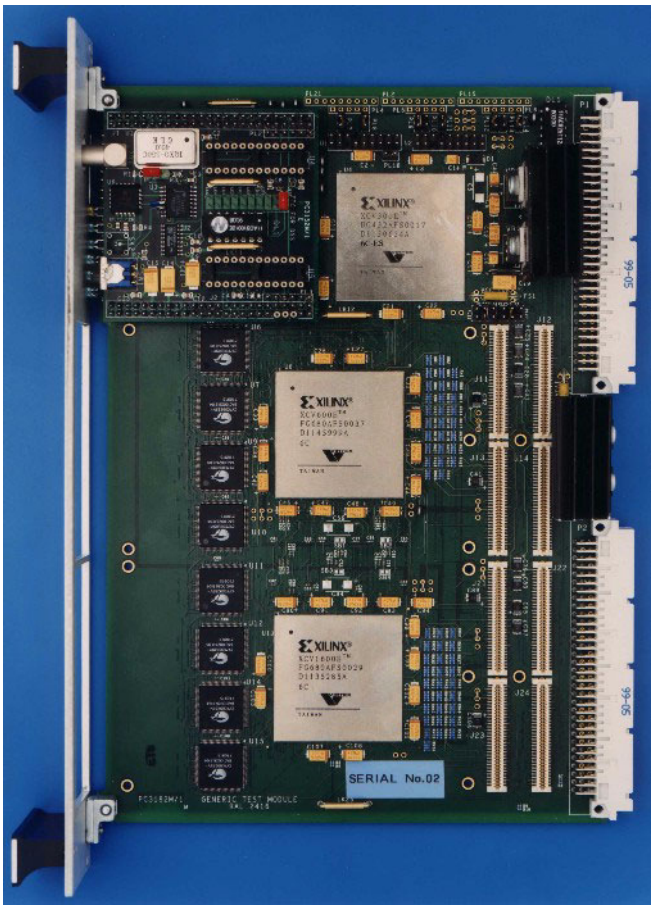

Figure 4.12 - The Generic Test Module

4.6.1 Test Setup

A more sophisticated top-level architecture was possible in the GTM. A diagram of the top-level design is shown in Figure 4.13. Interfaces are provided to VME, DPMs, the daughtercard and the other processor FPGA. Each bus is wide enough to act as the input or output of an algorithm. In the test described below, the VME and daughtercard buses were not used. Data was passed from one set of DPMs, through both FPGAs, then written to the other set of DPMs. A single signal from the control FPGA (accessible via VME) enables or disables the DPM read/write in the processor FPGAs, together with the inter-FPGA bus. Though both Pro1 and Pro2 can be used for data processing, only Pro1 was used in these tests. Pro2 simply passed the data straight through without altering it.

A VME controller running Linux was used to control the tests. Two C programs were written for different tests. The first (Program 1) was mainly used for debugging purposes. It reads data from a file, writes it to the Pro1 DPMs, toggles the ‘run-test’ signal briefly, then reads back the contents of all DPMs. The contents of the Pro2 DPMs are then compared with the expected output of the algorithm by eye. The second program (Program 2) was used to test the

full 24-to-4 sort. The program runs continuously, generating pseudo-random input data and predicting the output using a bit-level simulation of the sort algorithm. If the contents of the Pro2 DPMs differs from that prediction, the program alerts the user to the error.

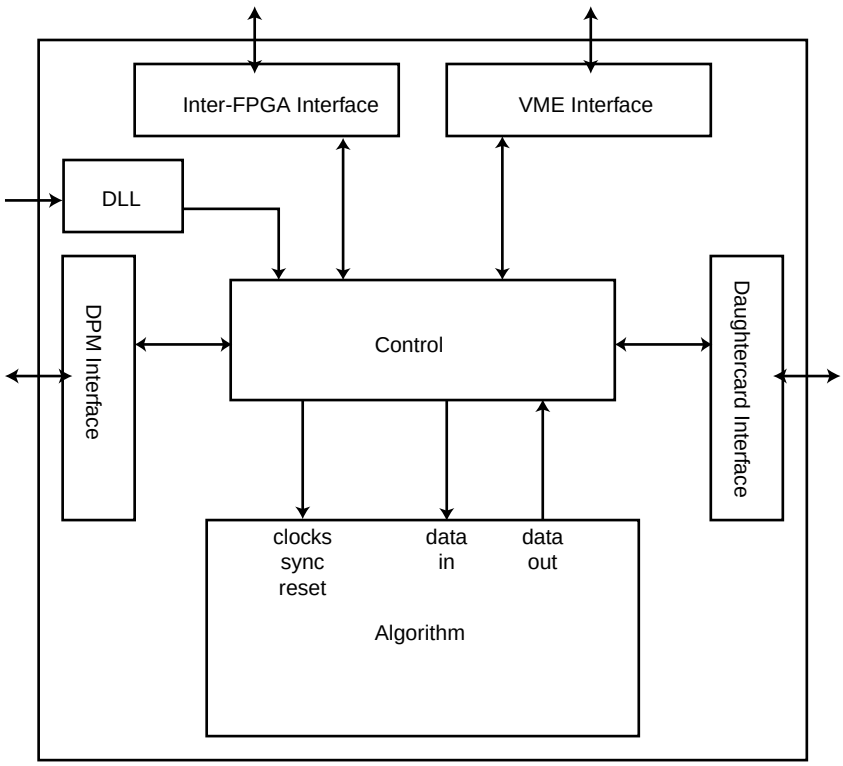

Figure 4.13 - Diagram of the top-level design used for algorithm tests in the GTM Pro1 FPGA

4.6.2 Pre-Sort

The pre-sort algorithm described in §4.3.2 was loaded into Pro1 using an 'Algo' entity as shown in Figure 4.14. De-multiplexers take the input data from 80 MHz down to 40 MHz, and multiplexers at the output reverse the process. Program 1 was used to compare predicted with real output data. No errors were observed in ~ 50 BX of data.

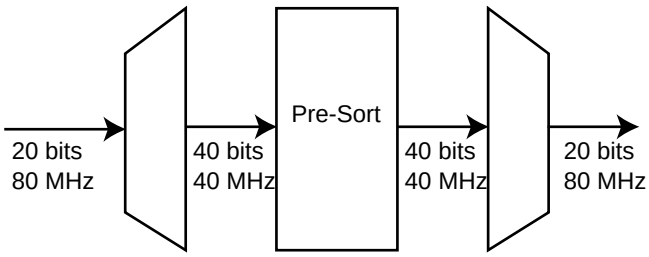

Figure 4.14 - 'Algo' entity used for testing the pre-sort.

4.6.3 Main Sort

The GTM was loaded with the main sort algorithm described in §4.3.3. As for the pre-sort algorithm, demultiplexers and multiplexers were used to process data at 40 MHz whilst communicating at 80 MHz. Again, Program 1 was used to compare real output with predictions, and again, no errors were observed in 50 BX of data.

4.6.4 Full '24 to 4' Sort

Finally, the GTM was loaded with a full sort processor using an algo entity as shown in Figure 4.15. This contains demultiplexers to reduce the input data rate to 40 MHz, a pre-sort for each input group, a main sort, and multiplexers on the output to transmit data at 80 MHz. This design used 19% of the logic available in the XCV600E. This figure is well within the effective maximum of 50-60%, above which it becomes excessively difficult to 'place and route' the design while maintaining the desired clock speed.

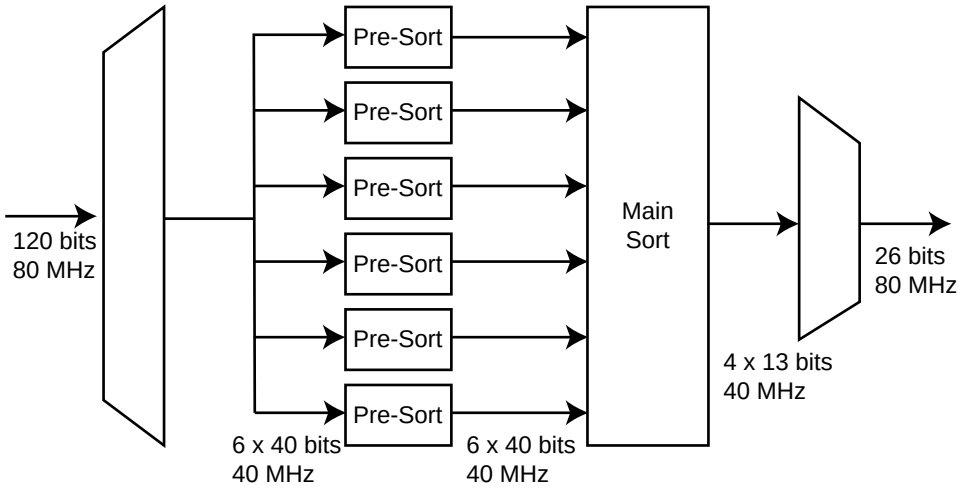

Figure 4.15 - 'Algo' entity for testing the full sort processor

Initial tests with Program 1 showed no errors, so Program 2 was used to test the algorithm over a long period. The test crate was left an a continuous cycle of writing input data to, and reading output data from, the GTM for approximately one week. Due to the relatively low data transfer rate over the VME backplane, and the need to reset the DPMs each cycle, this corresponded to ~ 6 minutes of data processing time, or $\sim 1.6 \times 10^{10}$ LHC bunch crossings. Since the error rate in the GCT is likely to be dominated by bit errors during transmission over cables, there seems little to be gained in extending the test period.

4.7 Conclusion

Gate-level functional simulation suggested that FPGA technology might be a suitable

processing solution for the Global Calorimeter Trigger. Initial hardware tests confirmed that it should be possible to implement a sort processor, operating at 80 MHz, in a Virtex FPGA. Finally, designs for a full sort processor were produced and tested using the Generic Test Module.

Chapter 5 - Performance of the Level-1 Jet and Missing Energy Triggers

5.1 Introduction

The primary goal of the LHC; discovery of the Higgs boson, will rely mostly on high p_t leptons to provide the Level-1 trigger signature. The jet trigger will also be used in some Higgs searches, notably those where the Higgs is produced in association with $t\bar{t}$. However, the jet trigger is most important in the search for Supersymmetry. In particular, an inclusive search of missing E_t plus multiple jet events currently provides our best hope of discovering or eliminating supersymmetry at the LHC [45]. Also in the context of Supersymmetry, the hadronic decays of τ leptons from a SUSY Higgs will require use of the jet trigger. The E_t^{miss} trigger is highly important in any channel involving the SUSY LSP, and is essential to capture any invisible decay modes of the Higgs.

The main obstacles to be overcome by both the jet and E_t^{miss} triggers are the effects of pile-up. The large number of relatively low energy particles resulting from pile-up interactions can combine to produce fake jets and missing E_t. The jet trigger also has to contend with an extremely high rate of real jets from QCD processes.

In this chapter, the performance of the Level-1 Jet and missing E_t algorithms is evaluated at the initial LHC luminosity of 2×10^6 $mb^{-1}s^{-1}$, by calculating the resolution, efficiency and rate of these triggers. This was achieved by processing a large dataset of QCD events with the CMSIM detector simulation program and ORCA, the CMS reconstruction program. This work follows on from previous studies [46], which used a less sophisticated technique for rate calculation, and which considered machine luminosities of 10^6 and 10^7 $mb^{-1}s^{-1}$. The authors of [46] have also undertaken similar studies [47], of which this work was independent.

5.2 Simulation Methods

5.2.1 Monte-Carlo Samples

The Monte-Carlo data samples used in the studies described below were generated by the CMS Monte-Carlo production team at the Fermilab National Accelerator Laboratory (FNAL), USA. The main type of sample used here consisted of QCD dijet events (the dominant process at the LHC). Since the $\hat{p}_t$ spectrum of these processes falls steeply ($\hat{p}_t$ being the transverse momentum exchanged by the struck partons), the events were generated in bins of $\hat{p}_t$, to ensure good statistics over the full spectrum.

The samples were generated using PYTHIA 6.152 [48], with MSEL=1 (all dominant inelastic processes) and cuts placed on $\hat{p}_t$ using CKIN(3) and CKIN(4).The $\hat{p}_t$ cuts for each sample follow from the previous sample (see Table 5.1), such that a continuous spectrum of $\hat{p}_t$ is obtained.

Full GEANT [49] detector simulation was carried out using CMSIM v120 [50], followed by digitization using ORCA 4.5.4 [51], including the addition of pile-up events. The full crossings were saved to a federated database on computers at FNAL, and further ORCA code was used for the analysis described below.

Sample name	$\hat{p}_t^{\,low}$ / GeV	$\hat{p}_t^{\,high}$ / GeV	N_{events}	σ/mb
jm_hlt1015	10	15	95,546	8.868
jm_hlt1520	15	20	129,450	1.854
jm_hlt2030	20	30	106,270	7.819×10^{-1}
jm_hlt3050	30	50	185,299	1.849×10^{-1}
jm_hlt5080	50	80	143,995	2.433×10^{-2}
jm_hlt80120	80	120	137,992	3.359×10^{-3}
jm_hlt120170	120	170	48,210	5.654×10^{-4}
jm_hlt170230	170	230	9,500	1.163×10^{-4}
jm_hlt230300	230	300	20,080	2.812×10^{-5}
jm_hlt300380	300	380	8,000	7.848×10^{-6}
jm_hlt380470	380	470	7,999	2.396×10^{-6}
jm_hlt470600	470	600	3,000	9.249×10^{-7}
jm_hlt600800	600	800	4,300	2.903×10^{-7}
jm_sm_qq_qqh120_inv	*n/a*	*n/a*	19,500	*n/a*

Table 5.1 - Monte-Carlo samples generated for trigger studies

In order to obtain good statistics for $E_t^{miss} > 100$ GeV in the resolution and efficiency calculations, a signals sample was used. The signal in question was that of an invisibly decaying Higgs boson produced via weak boson fusion. This sample was again produced using PYTHIA, with ISUB=123 and 124 for *H* production via *W* and *Z* fusion respectively. The invisible decays were produced by forcing the Higgs to decay to *ZZ**, which were subsequently forced to decay to neutrinos.

5.2.2 Pile-up

At the LHC, the luminosity will be sufficiently high that more than one *pp* interaction is obtained in each bunch-crossing. The mean number of events per bunch crossing, μ, is given by

$$\mu = \frac{L\sigma_{PU}}{R} \tag{5.1}$$

where L is the machine luminosity, R is the rate of filled bunch crossings (31.6 MHz) and σ_{PU} is the total cross-section for pile-up processes under consideration (the total inelastic *pp* cross-section at $\sqrt{s}$ = 14 TeV is expected to be 55.2 mb, calculated using PYTHIA). For the LHC design luminosity (10^7 $mb^{-1}s^{-1}$) this gives 17.5 events per crossing, while at the initial luminosity (2×10^6 $mb^{-1}s^{-1}$) we may expect 3.5. The actual number of events per crossing is of course distributed according to a Poisson distribution:

$$P_n = \frac{\mu^n}{n!}e^{-\mu} \tag{5.2}$$

The simulation of pile-up is straightforward. As mentioned in §3.11.1, the pile-up events are added at the 'hit' stage. A sample of unbinned minimum-bias events (i.e. MSEL=1 in PYTHIA, with no upper limit on $\hat{p}_t$ included via CKIN(4)) is used to provide pile-up. The combination of a primary event together with pile-up is referred to here as a *crossing*, whereas the term *event* refers to a single *pp* interaction.

5.2.3 Jet Finding and Calculation of E_t^{miss}

To provide a benchmark against which we can compare the Level-1 trigger results, various quantities are calculated using information available about the raw event at the 'particle' level. The input to these calculations are the stable particles available in the event record (the HEPEVT ntuple provided as input to CMSIM). Since the Level-1 algorithms deal only with calorimeter data, the particles that do not interact with the calorimeters are ignored. Here, these are taken to be muons, neutrinos and SUSY LSPs. Also, a fiducial cut of $|\eta| < 5$ is applied, corresponding to the pseudorapidity coverage of the calorimeters. It should be noted that the HEPEVT information was accessed using ORCA, and in the case of crossings that include pile-up, only the *primary event* was used. This is of importance when calculating rates (see §5.5.1 for further details).

Jets are found using the ORCA JetFinder package [52], specifically with an iterative cone algorithm. This algorithm forms a proto-jet by throwing a cone in $\eta-\phi$ around the direction of

the seed particle. The direction of the proto-jet is calculated from the energy weighted sum of its constituents, and another cone thrown around this new direction. The iteration stops when both the change in energy and position between iterations are below tunable thresholds. The thresholds used in this study were $\Delta E_t < 1\%$ and $\sqrt{\Delta\eta^2 + \Delta\phi^2} < 0.01$.

The process is repeated until all the particles in the event have been used. In practice, though, a threshold is placed on the seed energy (here, 1 GeV) to avoid wasting computation time on very low energy jets.

The E_t^{miss} calculation is given by the vector sum of particle momenta, again after application of a fiducial cut, and excluding unstable or non-interacting particles.

$$E_t^{miss} = \sqrt{E_x^2 + E_y^2} \tag{5.3}$$

where the *x*-component is given by:

$$E_x = \sum E_t \cos\phi \tag{5.4}$$

Simulation code for the Level-1 Calorimeter Trigger is currently contained in the 'L1CaloTrigger' subpackage of ORCA, the CMS reconstruction software. This code performs a bit-level functional simulation of the entire Level-1 calorimeter trigger. The data written out to ntuples for analysis consisted of E_t, η and ϕ for the 4 highest candidates in each jet stream (central, tau and forward), E_t^{miss} magnitude and ϕ-direction, and E_t^{total} (the scalar E_t sum over towers). This is essentially the data available to the Global Trigger. It should be noted that the purpose of the tau jet stream is to provide a trigger for hadronic decays of τ leptons, with an improved signal to background ratio over that of the 'standard' jet algorithm. However, since the tau and central jet streams are mutually exclusive, candidates for normal (i.e. non-τ) jet triggers should be taken from both.

Initial results were calculated from a small sample of the available crossings using an ORCA application written by the author. This application retrieves the generator level data from the federated database, calculates E_t^{miss} and uses the JetFinder library to find generator level jets. It then uses the L1CaloTrigger package to obtain the Level-1 Calorimeter Trigger output for the crossing. Resolution, efficiency and rate histograms are then filled with the relevant quantities and written out using the CHBook4 ORCA subsystem (a wrapper around the HBook package, allowing its use in C programs). Simple manipulation of the histograms was then performed with PAW, which was also used for graphical display of the results.

The computation time involved in running this code over the entire binned minimum bias dataset is substantial, and would duplicate other efforts, so the final plots shown here were calculated from HBook ntuples generated by the CMS Monte-Carlo production team at FNAL. The ORCA application used to produce these ntuples was written by the CMS Jet & Missing E_t Physics, Reconstruction and Selection (JetMEt PRS) group which uses identical calls to the ORCA application written by the author.

5.2.4 Minimum Bias Characteristics

The dominant process in a hadron collider such as the LHC is QCD jet production. The MSEL=1 switch (inelastic minimum bias) in PYTHIA simulates the following processes:

$$q_i\, q_j \rightarrow q_i\, q_j$$
$$q_i\, \overline{q}_i \rightarrow q_k\, \overline{q}_k$$
$$q_i\, \overline{q}_i \rightarrow g\, g$$
$$q_i\, g \rightarrow q_i\, g$$
$$g\, g \rightarrow q_k\, \overline{q}_k$$
$$g\, g \rightarrow g\, g$$

Note that, although no higher-order processes are included, multi-jet events are produced through initial- and final-state radiation.

Since these processes are the main background to jet triggers, the events are characterized before embarking on an evaluation of the trigger. At the very least, we can estimate the background rate of real jets - an indication of what would be possible given a perfect jet trigger. The distributions in this section were produced using the binned minimum-bias sample described above, and rate calculation method 1, discussed in §5.5.1.

Figure 5.1 shows the spectrum of $\hat{p}_t$, the transverse momentum exchanged by the struck partons. Clearly, this dictates the energy scale of the jets produced.

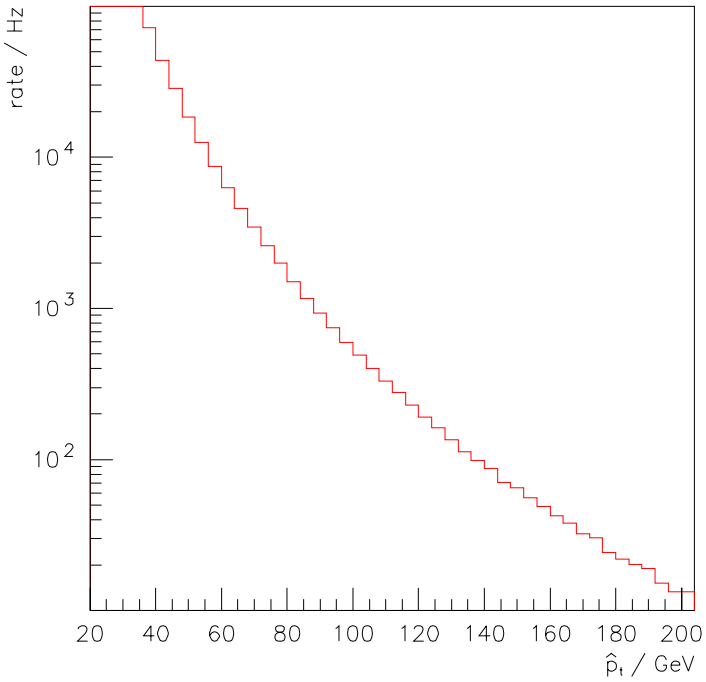

Figure 5.1 - $\hat{p}_t$ distribution in LHC minimum bias events

Figure 5.2A shows the jet E_t distribution. This roughly equivalent to the $\hat{p}_t$ distribution multiplied by a factor of two, because the events are di-jet events to first order. At this stage, it should be noted that when planning the Level-1 trigger thresholds, the total Level-1 accept rate is divided by a factor of three, as a safety margin to take into account the error on the QCD rate. The real total Level-1 accept rate is 100 kHz, so after taking into account the safety margin, a total rate of 33 kHz is assumed. This rate is divided between electron, muon, jet and combined triggers. The total rate available for jet triggers is likely to be no more than 10 kHz, say. As indicated on Figure 5.2A, this corresponds to a threshold on single jets at the generator level of ~65 GeV. If the Level-1 algorithm is assumed to be fully efficient, but susceptible to fake jets, this curve represents the lowest rate that is likely to be achieved. Since, in addition to this, a variety of multi-jet triggers will also be needed, the single jet threshold is likely to be above 100 GeV, which corresponds to a Level-1 accept rate of ~1 kHz.

Figure 5.2B shows the distribution of jets in pseudorapidity. It is clear from this plot that the very forward detectors (marked ‘HF’ in the range $3 < |\eta| < 5$) can expect a lower rate of jets by a factor ~3. However, a side effect of the high magnetic field necessary for high precision momentum measurement is that low energy particles deviate substantially by the time they reach the calorimeters. Particles with p_t below ~0.8 GeV do not reach the front face of the barrel ECAL, and will be swept forward to the endcap and forward calorimeters. Fortunately,

the very forward calorimeters used in CMS provide reasonably good rejection of low E_t particles.

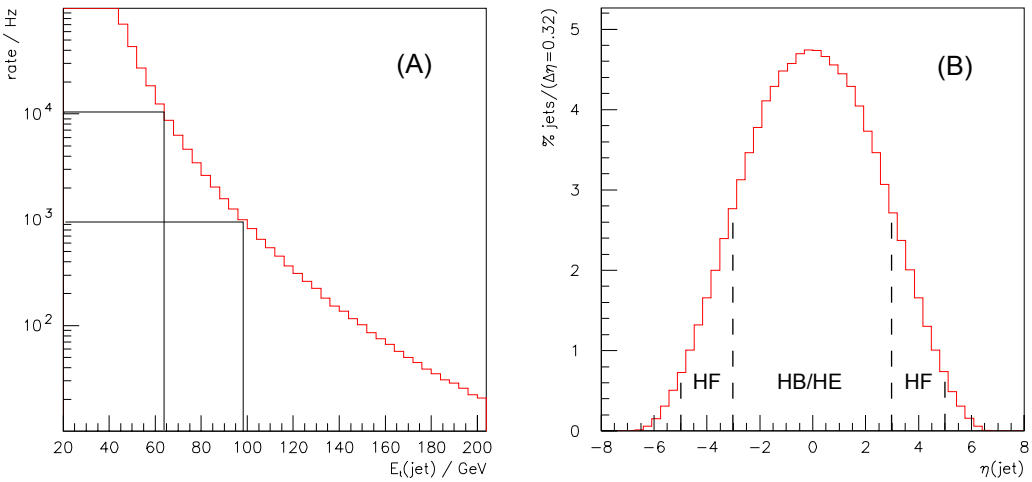

Figure 5.2 - Jet distributions from minimum bias events. The E_t thresholds corresponding to rates of 1 kHz and 10 kHz are marked, along with the calorimeter coverage in pseudorapidity.

The distribution of missing transverse energy from minimum bias events is shown in Figure 5.3. The rate falls steeply, but this plot does not account for the effects of pile-up. The rate of fake E_t^{miss} arising from overlapping events from different *pp* interactions is significant, and presents a major problem for measurement of missing E_t at the LHC.

To demonstrate how pile-up events can create large fake E_t^{miss}, consider the addition of two vectors, each of unit length, but with different directions. This is analogous to the E_t^{miss} vectors of two pile-up events. The magnitude of the vector sum is given by

$$\sqrt{2+2\cos(\Delta\phi)} \tag{5.5}$$

where $\Delta\phi$ is the difference in ϕ-direction between the two vectors. The root-mean-squared value, averaged over $\Delta\phi$, is $\sqrt{2}$. Although the mean value of E_t^{miss} from individual pile-up events is small, it should be clear that the addition of further events can lead to a significant rate of fake missing energy.

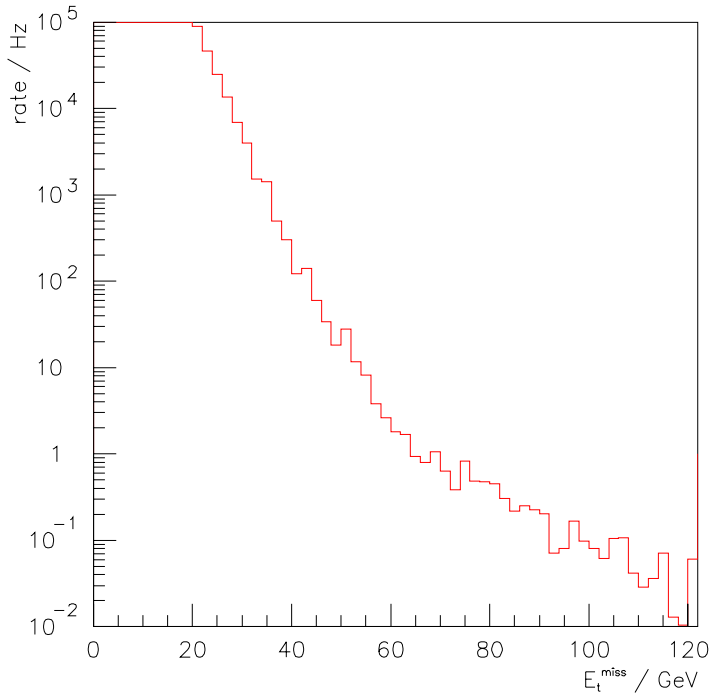

Figure 5.3 - E_t^{miss} distribution from minimum bias events (after fiducial cuts)

There is a clear change in slope of the E_t^{miss} distribution at ~60 GeV. The rate above this point is comparable to that before applying the fiducial cut of $|\eta| < 5$, while below this point the rate shows an increase after the cut [53]. Since the events under consideration are mostly dijet events, it seems likely that the increase in rate below 60 GeV is due to removal of one jet by the fiducial cut, but not the other. Above 60 GeV, this is less likely, because the rate of jets with sufficient E_t and $\eta > 5$ is extremely low.

5.2.5 Jet Matching

The algorithm described below was used to match particle level jets to their counterparts found by the Level-1 jet finder.

For each generator jet, working in order of descending E_t, we first find the nearest Level-1 jet in η-ϕ. If the separation between the two is less than some cut-off we match the jets. Above this cut-off the jets are unmatched. Figure 5.4 shows the distribution of the separation of each generator jet and the nearest Level-1 jet, for all generator jets with $E_t > 30$ GeV. The breadth of the distribution is due to the relatively coarse lateral granularity of the Level-1 trigger towers. For jets reconstructed in higher level triggers using the full granularity of the calorimeters, this distribution is narrower, and a lower cutoff can be used. The cut-off used in the following sections was $\Delta R < 0.5$.

The efficiency of the matching algorithm with this cut-off rises from ~25% for $10 < \hat{p}_t < 15$ GeV, to > 75% for $\hat{p}_t > 80$ GeV. The inefficiency arises both from the matching criteria, and the Level-1 jet algorithm. The poor efficiency for the lower $\hat{p}_t$ bins can be easily understood in terms of the latter. The Level-1 output is restricted to the highest 12 jets, from all events in the crossing. We match these jets against generator jets from the primary event only. In the lower $\hat{p}_t$ bins, jets from the primary event will have E_t comparable to those from pile-up events. Hence, the likelihood of the primary event jets reaching the Level-1 output is greatly reduced in comparison to the higher $\hat{p}_t$ bins, where the primary event jets have much greater E_t than those from pile-up events.

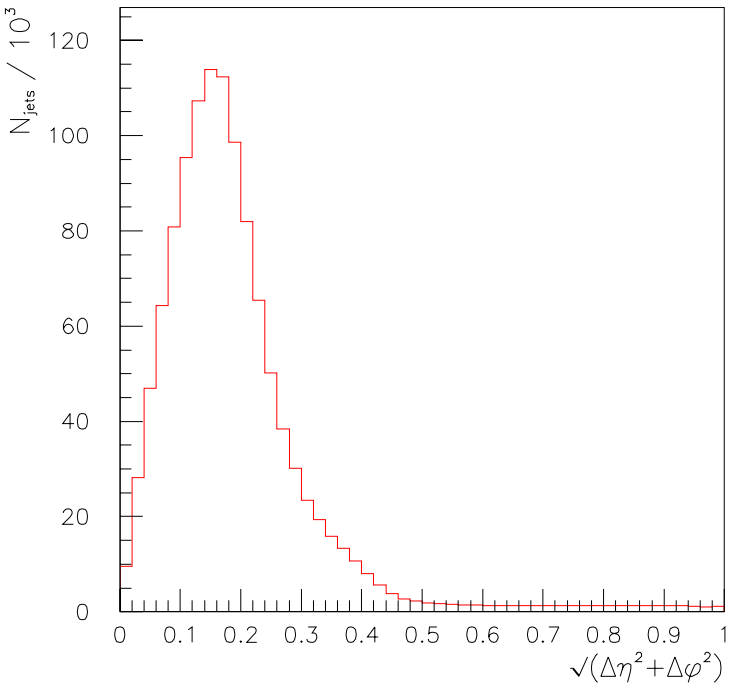

Figure 5.4 - Eta-phi separation of generator jet and nearest Level-1 jet

5.3 Jet E_t and E_t^{miss} resolution

The resolution of the jet transverse energy at Level-1 is calculated here, along with that of the missing transverse energy. Though the Level-1 trigger performance is judged in terms of trigger rate and efficiency rather than resolution, the rate is indirectly affected by the resolution. Since the rate of 95 GeV jets is higher than that of 100 GeV jets, if the jet resolution is 5%, then the trigger rate for a 100 GeV threshold may be near that of 95 GeV jets than that of 100 GeV jets. Before the jet energy resolution can be calculated, though, the energy scale must be calibrated.

5.3.1 Jet E_t Corrections

The final stage of the Level-1 jet algorithm is to apply a correction to the jet transverse energy. This correction is dependent on E_t and pseudorapidity, and is intended to counter the effects of non-linearity of the calorimeters, η-dependence of the jet response, and pile-up effects. The corrections are applied in hardware using look-up tables (LUTs), and may therefore take any form desired.

Simple linear and quadratic corrections were used in these studies, as described in [54]. The original study only calculated correction coefficients for luminosities of 10^6 and 10^7 $mb^{-1}s^{-1}$. The values of the coefficients used were calculated for luminosity of 2×10^6 $mb^{-1}s^{-1}$ using the method described in [54], by a member of the JetMEt PRS group [55].

In order to understand the nature of the corrections that must be applied to measured E_t, the raw response of the Level-1 jet finder is shown in Figure 5.5. Here, the ratio of measured to true E_t is plotted as a function of generator E_t and as a function of η. The response is clearly non-linear in both. The different response of the barrel + endcap and forward calorimeters can be clearly seen.

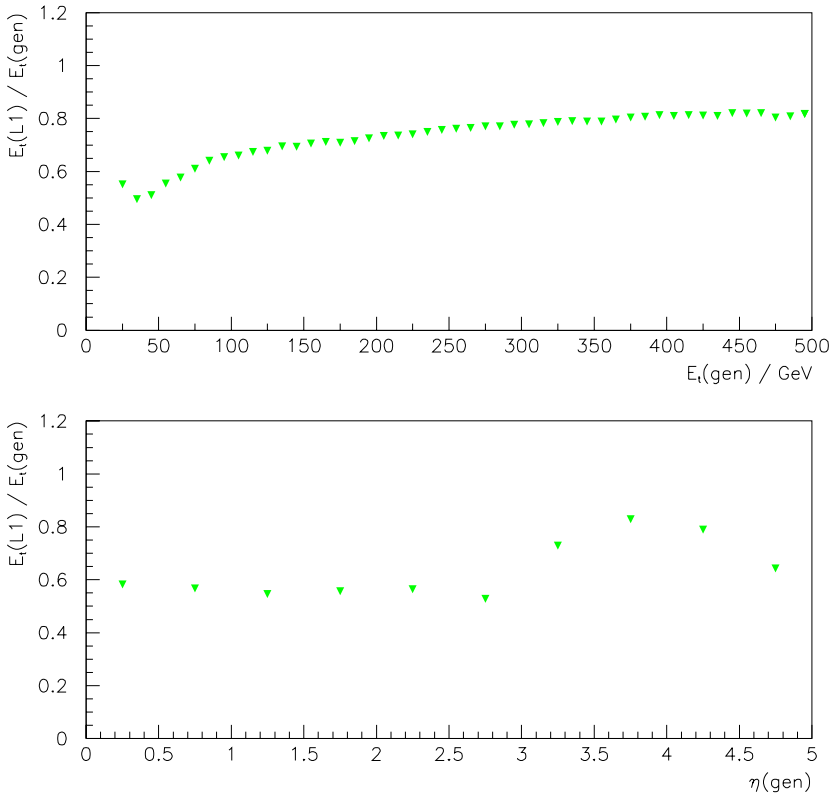

Figure 5.5 - Ratio of raw Level-1 jet E_t to true jet E_t, plotted as a function of both E_t and pseudorapidity

The correlation between measured E_t and generated E_t is plotted in Figure 5.6 for $0.0 < \eta < 0.2$ (corresponding to a single ring of trigger towers). The relationship is well described by a linear fit. However, the effect of pile-up at low energy distorts this relationship slightly, and the value of χ^2 can be improved slightly by including a small quadratic term:

$$E_t(L1) = p_1 E_t(gen)^2 + p_2 E_t(gen) + p_3 \quad (5.6)$$

However, the forward calorimeters offer improved rejection of pile-up, so a linear fit is used for jets in these regions (i.e. for $|\eta| > 3.0$).

The correlation between corrected E_t and generated E_t is also plotted in Figure 5.6. As can be seen, on average, the corrected transverse energy is in good agreement with the generated jet energy up to ~300 GeV.

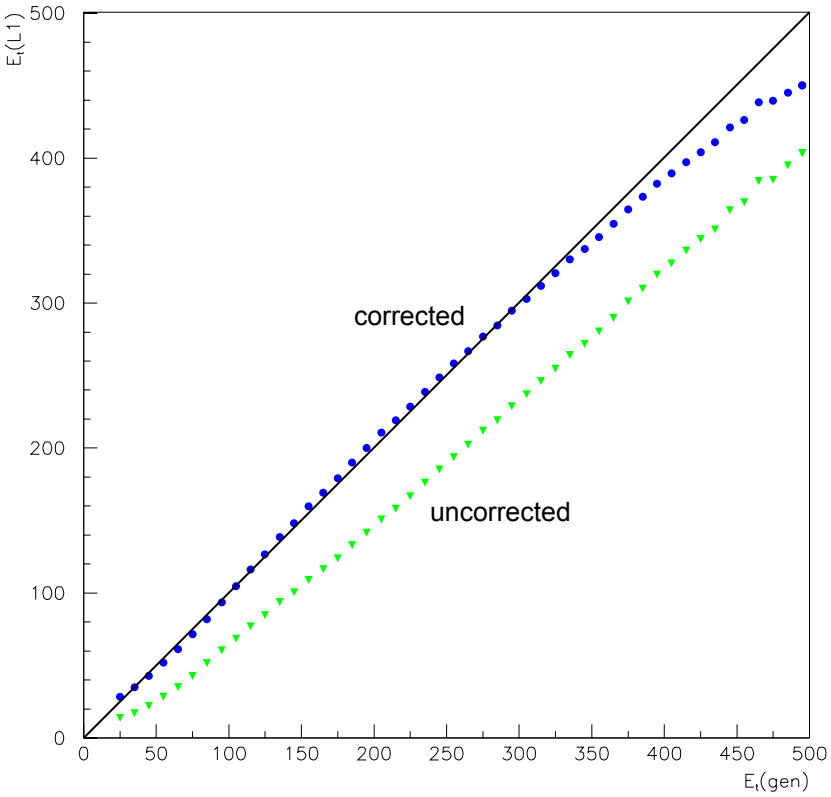

Figure 5.6 - Mean Level-1 jet E_t as a function of generator level E_t before (light lower curve) and after corrections (dark upper curve). The solid line is E_t(L1) = E_t(gen).

As explained in the previous section, we must apply a correction to the Level-1 jet E_t in order to recover the true E_t. However, there remains a spread in the distribution of this reconstructed E_t. The E_t resolution of the Level-1 trigger dictates the efficiency of the selection in the region of the threshold. If the spread in reconstructed E_t for a given true E_t is small, then the trigger will produce high efficiency for jets with true E_t slightly above the threshold. However, if the spread is large, the trigger will only produce high efficiency for jets with E_t well above the threshold.

5.3.2 Resolution Calculation

It is usual to parameterize the energy resolution of a calorimeter according to Eq. 5.7. The stochastic term, p_1, is due to sampling fluctuations, while the constant term, p_2, is due to intrinsic effects such as a differing response to hadrons and electrons.

$$\frac{\sigma_E}{E} = \frac{p_1}{\sqrt{E}} \oplus p_2 \oplus \frac{p_3}{E} \tag{5.7}$$

In this section we consider transverse energy rather than energy, but the form of parameterization remains as above.

Since we have access to the 'truth' after matching generator level jets to their Level-1 counterparts, we can find the fractional error for each jet energy measurement. The fractional resolution is then given by the root-mean-squared value of

$$E_t(L1)/E_t(gen) \tag{5.8}$$

Since the missing energy measurement depends on energy measurements made over the entire calorimeter, it is usual to parameterize the missing E_t resolution in terms of the total deposited transverse energy, E_t^{total}:

$$\frac{\sigma_{Etmiss}}{E_t^{total}} = \frac{p_1}{\sqrt{E_t^{total}}} \tag{5.9}$$

5.3.3 Results

Figure 5.7 shows the root-mean-squared value of $E_t(L1)/E_t(gen)$ (i.e. the fractional E_t resolution) as a function of E_t, for all Level-1 jets. Figure 5.8 shows the resolution for central and forward jets separately. Note that for this calculation the central and tau streams have been combined. The results of a fit to Eq. 5.7 are shown on each plot, and summarized in Table 5.2.

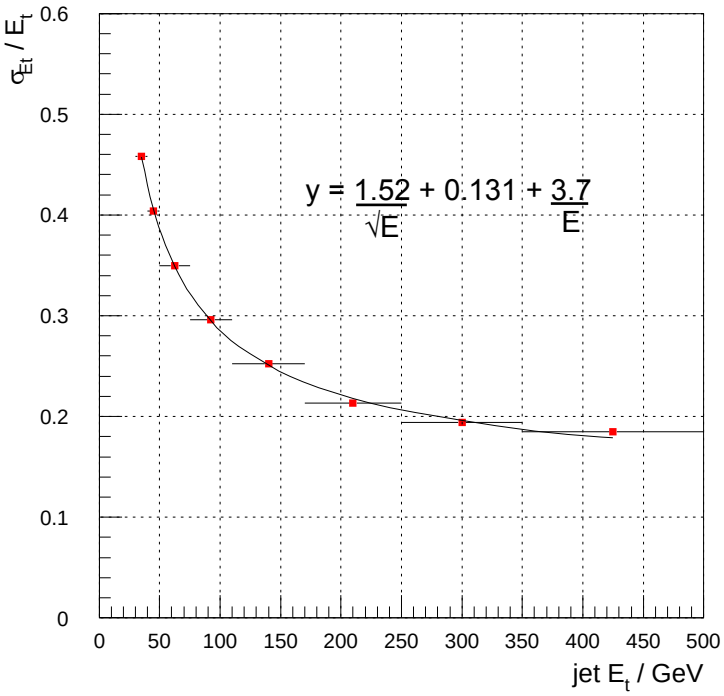

Figure 5.7 - E_t resolution for all Level-1 jets

Figure 5.8 shows the E_t resolution for central and forward jets separately. For the purposes of

this plot the central and tau streams are added together to produce Figure 5.8A. As can be seen, the central and forward jet resolutions are very similar below 150 GeV Et. Above this point the central jets appear to have the edge, although the statistics in the forward jet sample soon run out, so direct comparison is not possible.

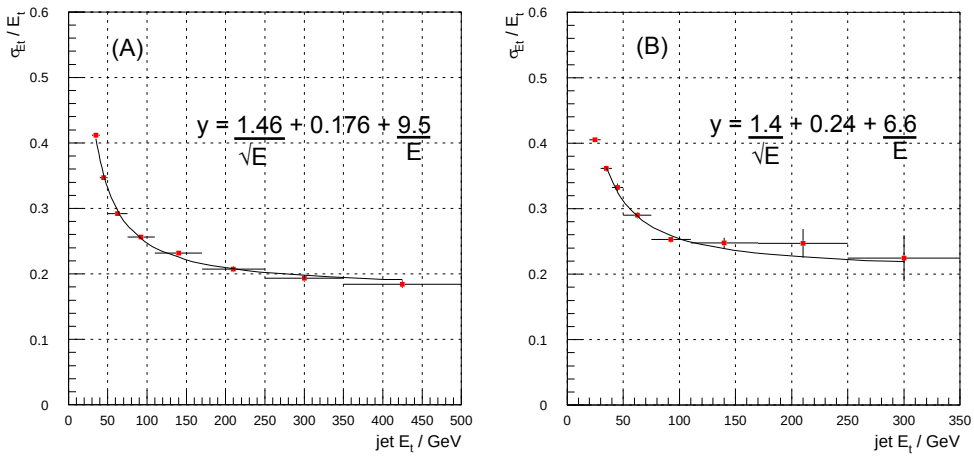

Figure 5.8 - E_t resolution for central / tau and forward jets

Jet type	p_1 / GeV$^{1/2}$	p_2	p_3 / GeV
All jets	1.52 ± 0.04	0.131 ± 0.004	3.7 ± 0.9
Central/tau jets	1.46 ± 0.10	0.176 ± 0.005	9.5 ± 0.6
Forward jets	1.4 ± 0.4	0.24 ± 0.02	6.6 ± 2.1

Table 5.2 - Jet resolution fit results

Figure 5.9 shows the resolution of the Level-1 missing E_t measurement as a function of E_t^{total}, obtained using the binned minimum bias sample. The result of a fit to Eq. 5.9 is

$$\frac{\sigma_{Etmiss}}{E_t^{total}} = \frac{0.922 \pm 0.002}{\sqrt{E_t^{total}}} \qquad (5.10)$$

It should be noted that the resolution obtained here is almost as good as that obtained from the vector sum of offline towers ($\sim 0.55\sqrt{E_t^{total}}$ [56]), although the offline measurement can be substantially improved by correcting for the hadronic component of the missing E_t [57].

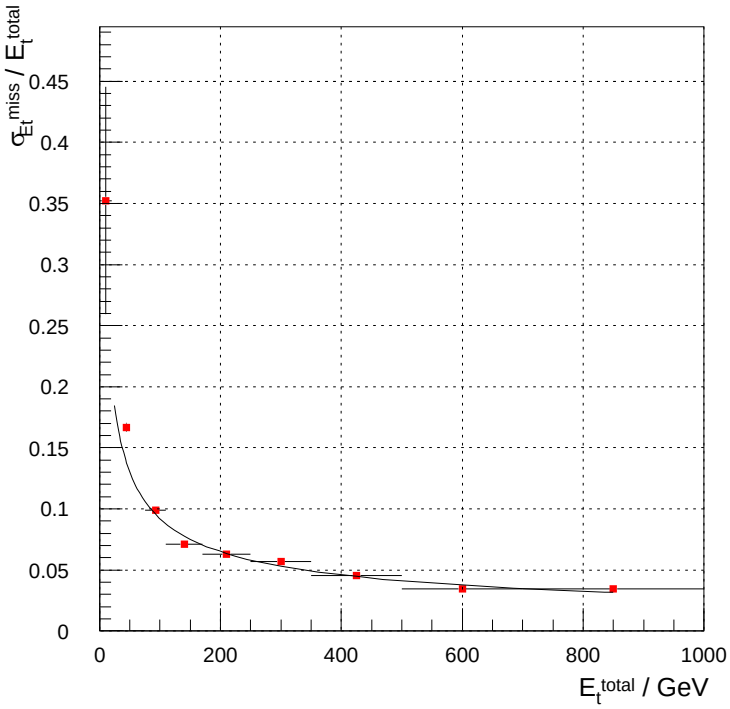

Figure 5.9 - Level-1 E_t^{miss} resolution

5.4 Trigger Efficiency

Since the trigger efficiency for a particular signal will depend on the threshold(s) used, and these in turn depend on the bandwidth available to the DAQ, a more useful measure at this level is the 'turn-on curve'. The efficiency in selecting jets (rather than events) is plotted against the true jet E_t, for a given Level-1 threshold.

5.4.1 Jet Trigger Turn-on

Figure 5.10 shows the efficiency of the Level-1 jet algorithm as a function of generator jet E_t for several Level-1 E_t thresholds. The efficiency is shown both before and after corrections (in light and dark grey, respectively). The generator jets were matched to Level-1 jets using the procedure described in §5.2.5. The value of generator jet E_t for which the trigger achieves 95% efficiency is marked on each plot. The improvement obtained by applying the corrections is quite clear.

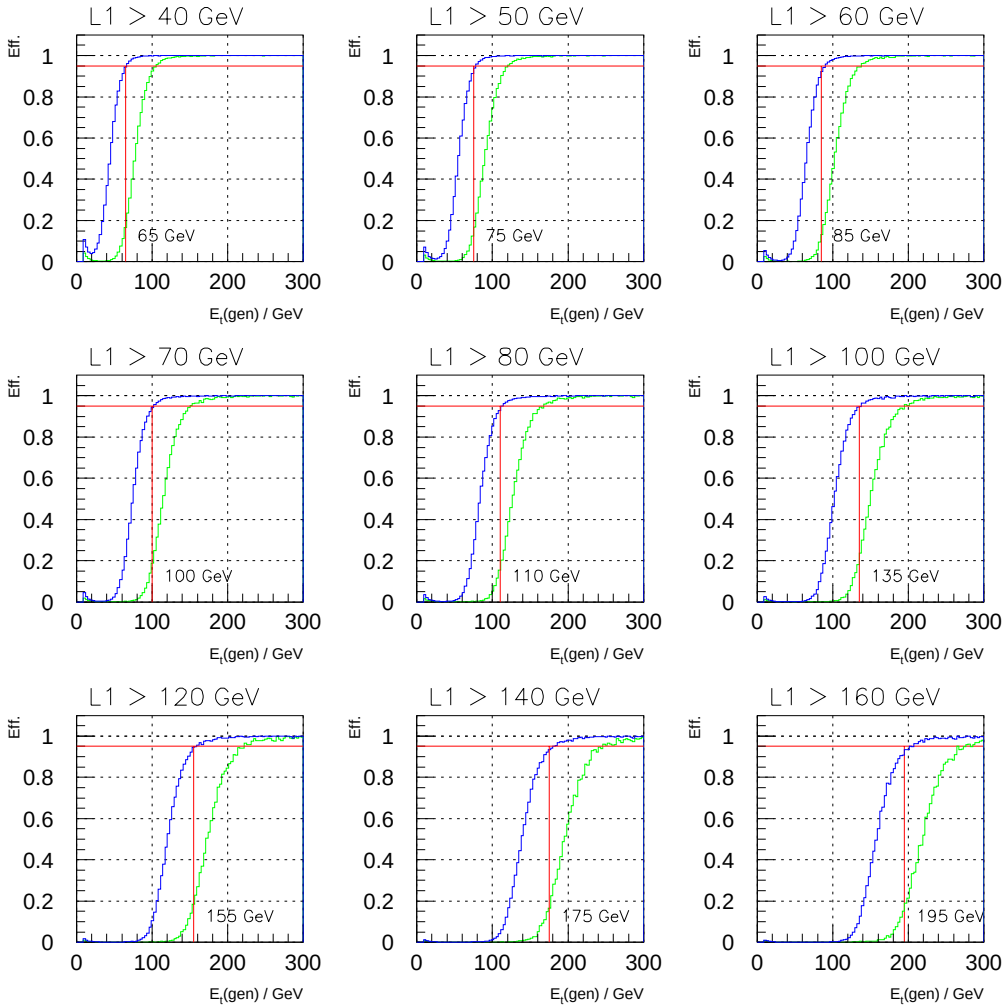

Figure 5.10 - Turn-on curves for Level-1 central / tau jets, before (light grey curve) and after (dark grey curve) corrections. 95% efficient points are marked with a solid line.

The turn-on curves for central and tau jets are much the same as those obtained using all jets, however, the forward jets offer a noticeable, though slight, improvement over the general case, as shown in Figure 5.11.

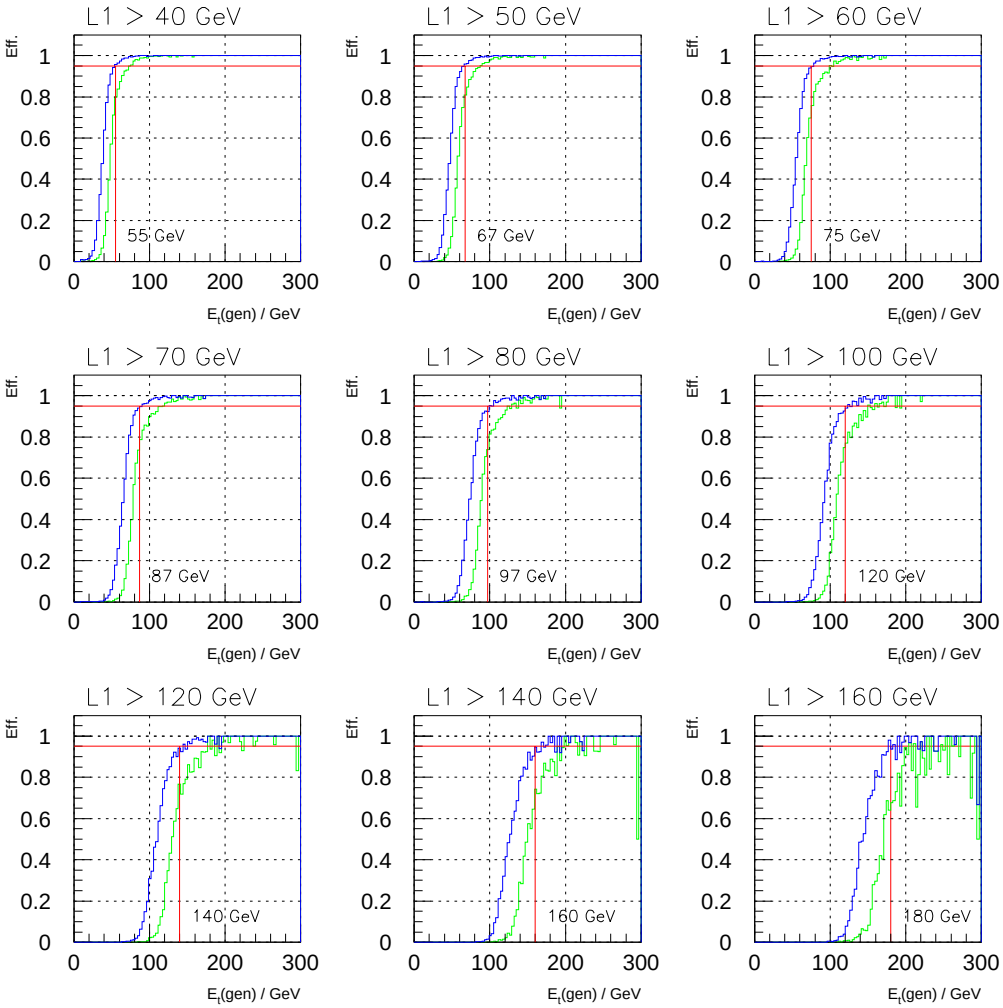

Figure 5.11 - Turn-on curves for Level-1 forward jets, before (light grey curve) and after (dark grey curve) corrections. 95% efficient points are marked with a solid line.

5.4.2 E_t^{miss} Trigger Turn-on

The turnon curve for the Level-1 E_t^{miss} trigger is shown in Figure 5.12. The efficiency is again plotted as a function of generator level E_t^{miss}. These curves are comparable to the jet turn-on before corrections are applied. This is expected, as the Level-1 missing E_t calculation does not cater for the non-linear response of the calorimeters.

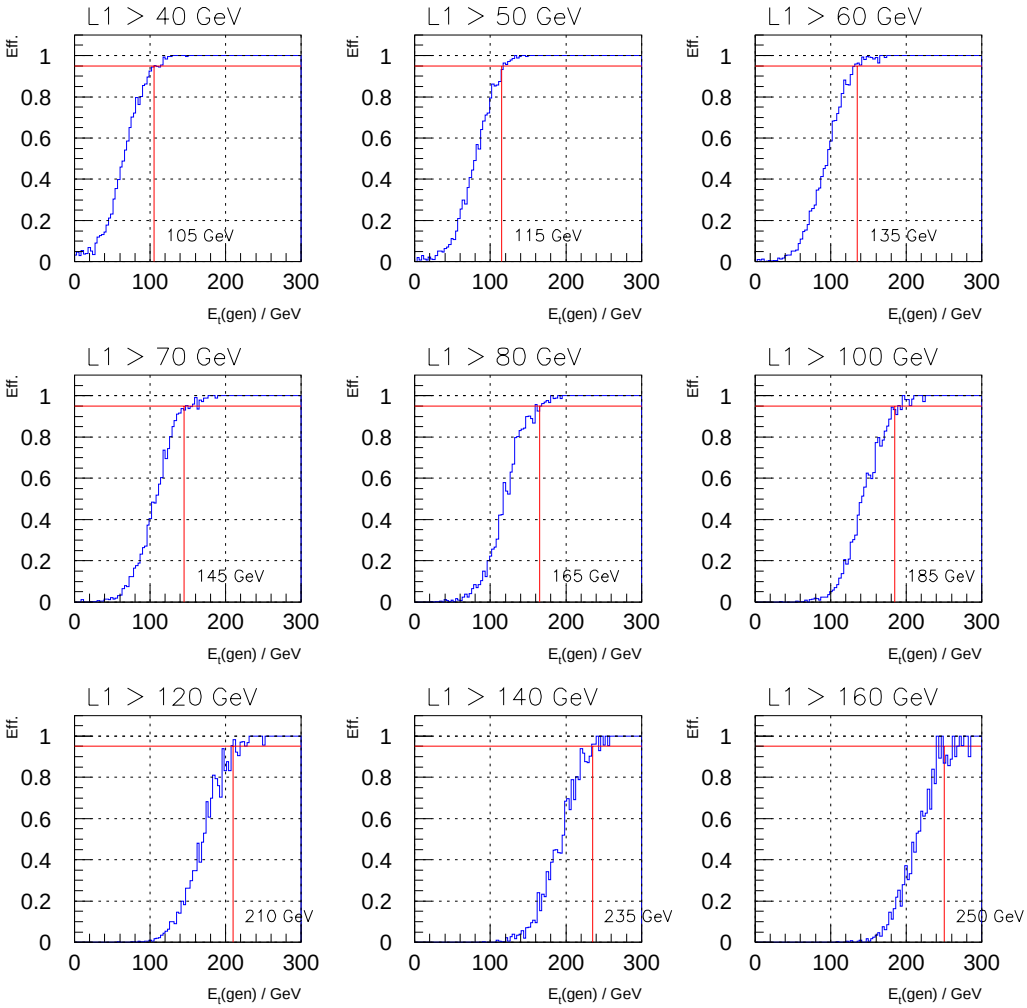

Figure 5.12 - E_t^{miss} algorithm turn on curves. 95% efficient points are marked with a solid line.

5.4.3 Parameterisation of Results

As a final exercise, the 95% efficiency points can be plotted against the corresponding Level-1 threshold. Figure 5.13A shows jets in the central region in red and jets in the forward region in blue. Figure 5.13B shows missing transverse energy. The relationship is clearly linear for all 3 sets of data. Results of fitting Eq. 5.11 are given in Table 5.3.

$$E_t^{95} = a \cdot E_t^{L1} + b \qquad (5.11)$$

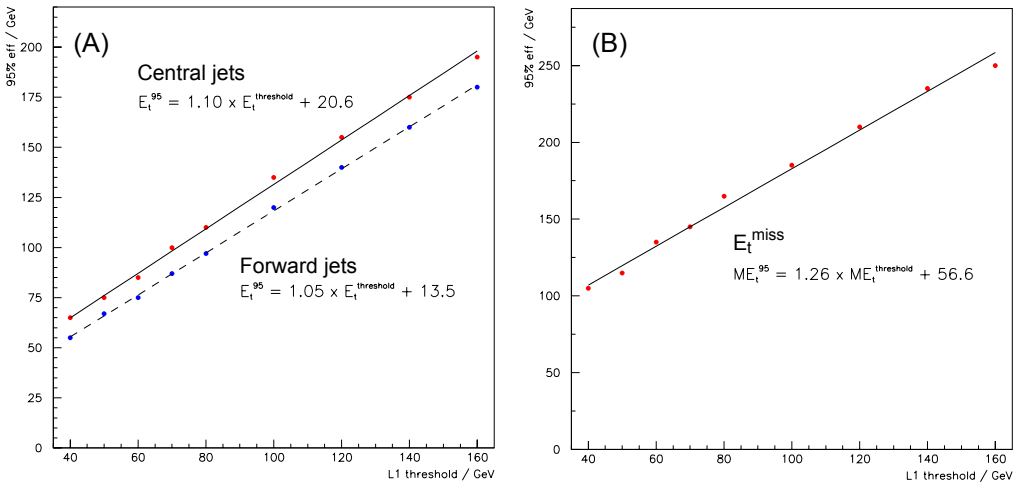

Figure 5.13 - 95% efficiency point plotted as a function of Level-1 threshold for A) central/tau jets & forward jets and B) E_t^{miss}.

Trigger Object	*a*	*b* / GeV
Central/Tau jets	1.10	20.6
Forward jets	1.05	13.5
E_t^{miss}	1.26	56.6

Table 5.3 - Results of fits to 95% efficiency vs. Level-1 threshold

5.5 Trigger Rate

Level-1 accept rates are shown here for various jet and E_t^{miss} triggers. The triggers under consideration are all constructed by applying E_t and η thresholds to the 12 jet objects (four each of central, forward and central τ-jet) and E_t^{miss} available at the Level-1 Global Trigger.

5.5.1 Rate Calculations

Estimation of trigger rates with a reasonable degree of confidence is clearly of importance, as the rate dictates the efficiency available for a given DAQ bandwidth. However, the presence of multiple interactions during each LHC bunch-crossing complicates this calculation.

Method 1

Given at most a single interaction per bunch crossing, the trigger rate can be calculated by producing a sample of events and counting those that pass the trigger. In practice it is usual to produce several samples, binned by the $\hat{p}_t$ of the interaction, to extend the $\hat{p}_t$ region with good statistics. The trigger rate for luminosity L is then given by

$$R = L\sum \sigma_i p_i \qquad (5.12)$$

where the sum runs over bins, σ_i is the cross-section of the *i*th bin and p_i is the proportion of events from the *i*th bin that pass the trigger.

Note that this method is also valid for calculating rates from samples that include pile-up, as long as the event that fires the trigger is the 'primary' event. This includes the 'generator level trigger' rates shown later.

Method 2

However, as already stated, the Level-1 trigger output is the result of the superposition of several events. The crossings used to calculate trigger rates consist of a single primary event from a binned minimum-bias sample, together with several 'pile-up' events from the unbinned minimum-bias sample. This complicates the rate calculation somewhat, since the pile-up may cause the trigger.

The problem is most easily seen if we consider a trigger caused by single objects originating from a single interaction, e.g. the single jet trigger. For high E_t thresholds, equation (5.12) remains a reasonable approximation, since the jet will almost always originate from the 'signal' event. If the trigger jet originates from a pile-up event, however, equation (5.12) attributes it the same weight as if it originated from the 'signal' event. In the lower $\hat{p}_t$ bins, we find that rare pile-up events with high $\hat{p}_t$ may cause a trigger and are given the weight of a lower bin. This results in a rate plot as shown in Figure 5.14. The vertical lines indicate the most extreme examples of this mis-weighting, where the jet (or missing E_t) causing the trigger is from an event with $\hat{p}_t$ well above that of the $\hat{p}_t$ bin used to weight it.

The correct method for weighting events is derived in [58]. The result is applicable to the Monte-Carlo method described above. To calculate the weight for a given crossing, we first find the distribution of $\hat{p}_t$ within that crossing. The weight is then given by

$$W(\vec{n}) = L\sum_i N_i \frac{n_i}{f_i} \tag{5.13}$$

where i runs over $\hat{p}_t$ bins; the ith component of $\vec{n}$, n_i, is the number of events in the crossing that fall into the ith $\hat{p}_t$ bin; and f_i is the fractional cross-section of that bin (i.e. σ_i/σ_T). The rate for a given trigger is then found by summing the weights of those crossings that pass the trigger.

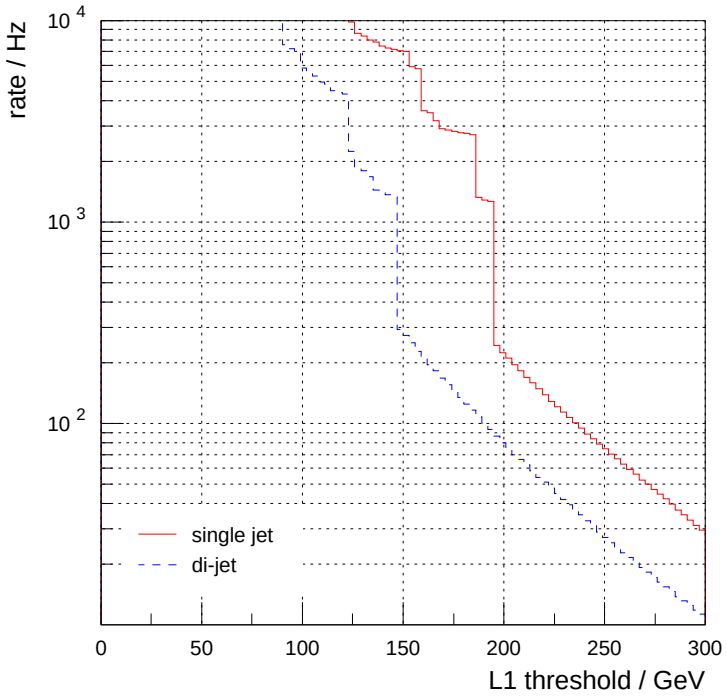

Figure 5.14 - Level-1 single and di-jet rate incorrectly calculated using Method 1.

5.5.2 Level-1 Rates

Figure 5.16 shows the trigger rate for single, di-, tri- and quad-jet triggers using all the jets available at the Level-1 output. The Level-1 trigger rate is shown with a series of points, while the corresponding generator level rate is shown as a solid line.

Since the Level-1 energy scale has been corrected, we may directly compare the generator level rate with the Level-1 rate. The latter is expected to be higher than the former, because of the breadth of the jet E_t resolution. For a given Level-1 threshold, some jets with generator E_t below the threshold will be mis-measured and fire the trigger. Since the jet spectrum falls

steeply with E_t, the Level-1 rate is dominated by jets with real E_t below the Level-1 threshold.

Since the Level-1 trigger rate is the result of the convolution of the generator level trigger rate with the Level-1 jet E_t resolution function (which is well represented by a Gaussian), we may estimate the expected Level-1 rate from these quantities. The E_t resolution was found earlier to be ~ 12% at 100 GeV, so we may assume the Level-1 trigger rate for this threshold is dominated by jets with true E_t of ~90 GeV. The generator rate for a 90 GeV threshold is 3.3 kHz, which is slightly lower than the Level-1 rate at 100GeV, of 3.6 kHz. The small difference may be attributed to the approximate method used.

The E_t^{miss} rate is plotted in Figure 5.17. It exhibits the same 'kink' that can be seen in the generator level distribution, albeit shifted to ~100 GeV. Clearly, the rate of fake E_t^{miss} is high. Insufficient statistics were available to calculate the rate beyond ~ 200 GeV, and a bug in CMSIM caused problems above this E_t, when a non-physical energy deposit is recorded in a calorimeter cell. These non-physical energy deposits are often substantially larger than the $\hat{p}_t$ of the event, and cause steps in the rate plot similar to those caused by the mis-weighting, but of lesser severity.

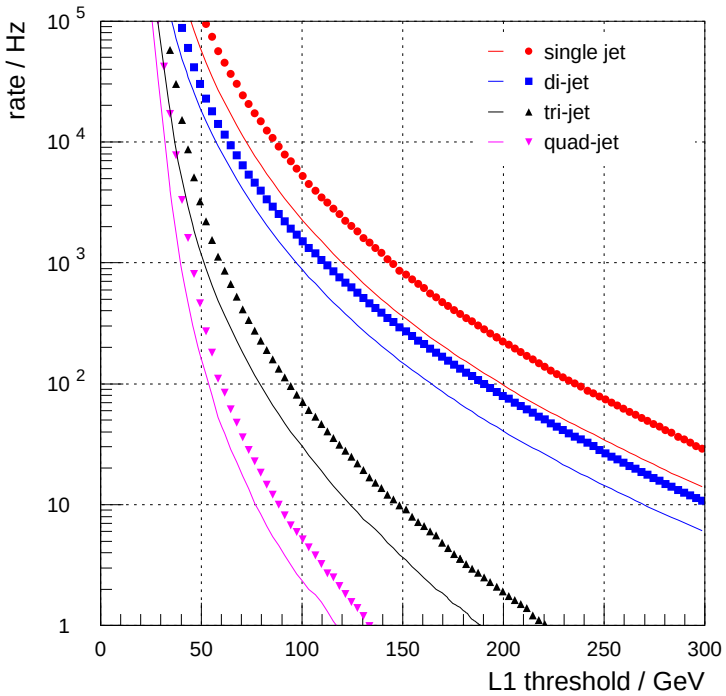

Figure 5.15 - Rates for 1, 2, 3 & 4 jet triggers

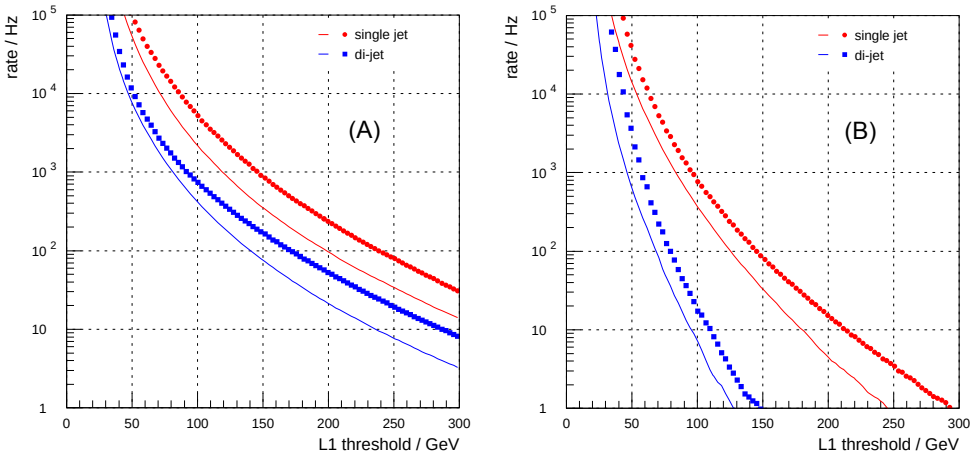

Figure 5.16 - Single jet and di-jet trigger rates, for A) central / tau and B) forward streams

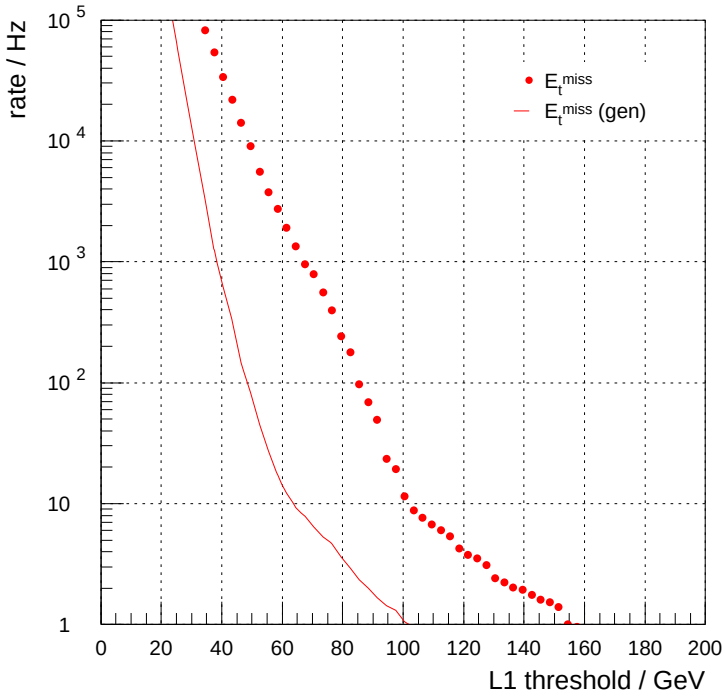

Figure 5.17 - E_t^{miss} trigger rate as a function of the Level-1 threshold

An alternative presentation of trigger rates is to plot the rate as a function of the 95% efficient point for that threshold. This is done by applying Eq. 5.11 to the abscissae of Figure 5.16 and Figure 5.18. The resulting plots are shown in Figure 5.18 and Figure 5.19, respectively. This is clearly presents the trigger in harsher terms, but can be useful when considering triggers for a particular set of offline cuts. One should not expect, for example, to collect single jet events

below the 100-150 GeV region with full efficiency, as the bandwidth required is unlikely to be available. The need for triggers that combine several objects in order to reduce the thresholds becomes clear when viewing these plots. This should not be seen as poor performance of the Level-1 jet trigger, however, but as one of the obstacles to physics analysis presented by a high luminosity hadron collider such as the LHC.

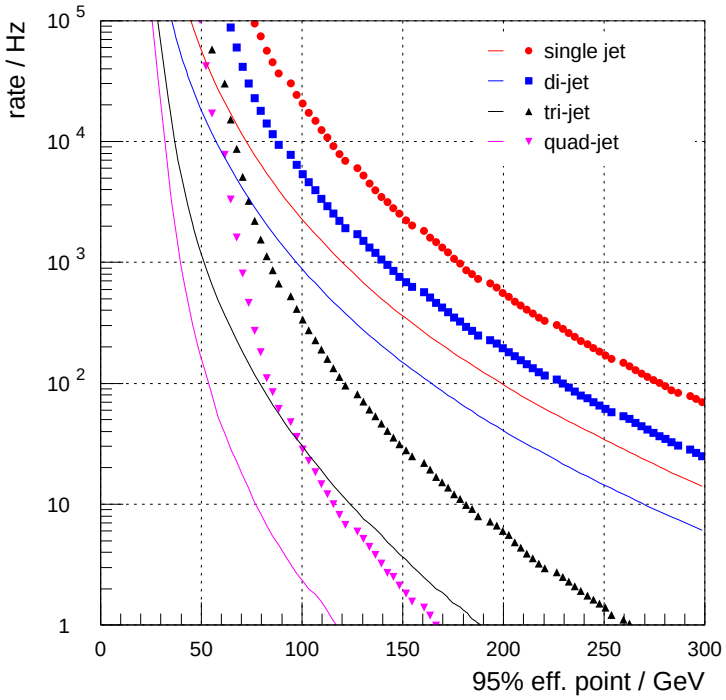

Figure 5.18 - Rates for 1, 2, 3 & 4 jet triggers, as a function of 95% efficiency point.

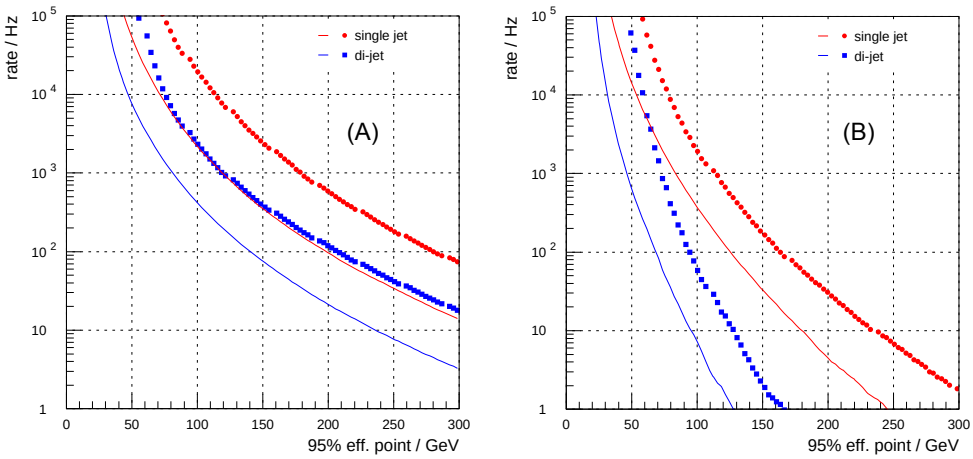

Figure 5.19 - Single and di-jet trigger rates for A) central / tau and B) forward streams as a function of 95% efficiency point.

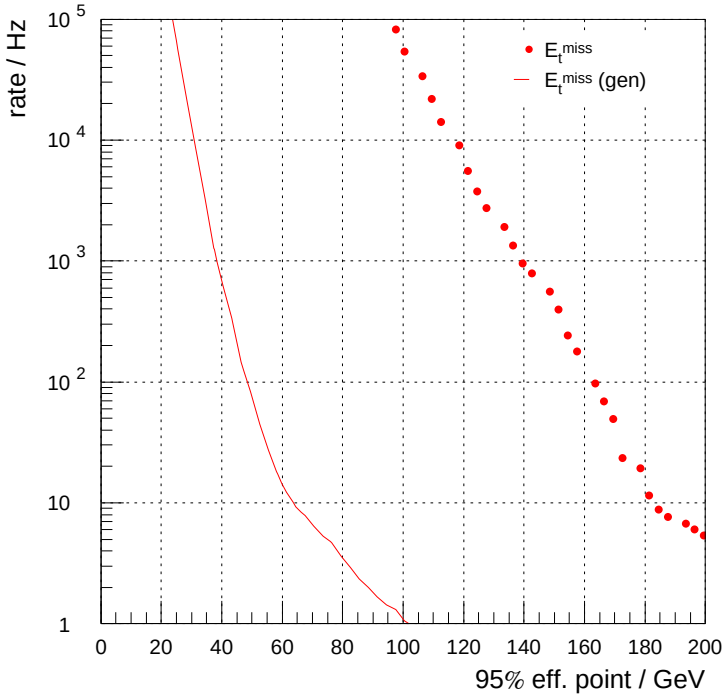

Figure 5.20 - E_t^{miss} trigger rate at low luminosity as a function of 95% efficiency point.

5.5.3 Low Luminosity Trigger Table

The allocation of bandwidth to different types of trigger can proceed in various ways. Various possible tables are assessed in [47], corresponding to different rate allocations to electron, jet and combined triggers. The final, most performant version is given in Table 5.4. In addition, the muon trigger thresholds are included, taken from [59]. This table is missing combined triggers, which have not yet been studied in sufficient detail at low luminosity and form part of the subject of Chapter 7.

This table is designed to maximize the Higgs physics selection, whilst maintaining good efficiency for other channels, such as W, Z and top signals. The total Level-1 accept rate is taken to be 50 kHz. This is below the design rate, since it is likely that only a portion of the full DAQ hardware will available at start-up. As mentioned earlier, a safety margin of a factor of three is allowed for the uncertainty in the rate calculation. This leaves ~16 kHz for the total calculated rate. Of this, 4 kHz each are allowed for electron and muon triggers, 7 kHz for jet triggers and 1 kHz for combined triggers. This allocation of bandwidth was found to provide better efficiency across the range of Higgs signals than simply dividing the rate equally between the four categories [47].

Trigger	Thresh. (GeV)	95% eff. (GeV)	Ind. Rate (kHz)	Cum. Rate (kHz)
e	21	27	3.9	3.9
ee	15	19	0.2	4.0
τ	85	–	4.9	8.8
ττ	75	–	0.7	8.8
j	110	134	3.2	10.4
jj	90	113	2.1	10.6
jjj	60	71	0.8	10.8
jjjj	50	53	0.3	10.9
e · j	10, 100	15, 125	0.4	11.0
e · τ	10, 75	–	0.8	11.2
E_t^{miss}	105	200	0.01	11.2
e · E_t^{miss}	10, 50	15, 140	0.4	11.5
j · E_t^{miss}	60, 60	80, 150	0.7	11.7
Total E_t	600	1200	0.04	11.7
H_t	400	470	0.6	11.8
e (non-isolated)	45	51	0.2	11.8
ee (non-isolated)	25	37	0.3	11.8
μ	14		3.5	15.3
μμ	5			15.8
Total				15.8

Table 5.4 - Official Level-1 trigger table, taken from [47]. Muon thresholds are taken from [59]. For the purposes of rate calculation, the muon and calorimeter triggers are assumed to contain no overlap.

It should be noted that the rate allocated to the missing energy trigger is very low (0.01 kHz).

5.6 Summary of Results

The performance of the CMS Level-1 jet and missing energy triggers have been quantified in several ways.

The jet E_t resolution was calculated, and the fit to a parameterisation including stochastic, constant and noise terms gives:

$$\frac{\sigma_{Et}}{E_t} = \frac{2.52}{\sqrt{E_t}} \oplus 0.131 \oplus \frac{3.7}{E_t} \tag{5.14}$$

which gives a resolution of 28.6 GeV for jets with transverse energy of 100 GeV.

The E_t^{miss} resolution was also calculated. Parameterized in terms of total transverse energy deposited during the bunch-crossing (E_t^{total}), the resolution is given by:

$$\frac{\sigma_{Etmiss}}{E_t^{total}} = \frac{0.922}{\sqrt{E_t^{total}}} \tag{5.15}$$

which gives a resolution of 29 GeV for events with 1 TeV of total transverse energy.

The jet trigger turn-on reflects the resolution figures; as a general rule for Level-1 thresholds below 200 GeV, the trigger reaches 95% efficiency within ~35 GeV (20 GeV) of the threshold for jets in the barrel/endcap (forward) calorimeters.

Likewise, the missing energy trigger turn-on reflects the corresponding resolution result. In general, for thresholds below 200 GeV, the trigger is 95% efficient within 90 GeV of the Level-1 threshold.

Finally, the trigger rate was calculated. For jets, it was found to be consistent with the rate of ‘generator level triggers’ given the resolution calculated above. For E_t^{miss} the rate of fake triggers is substantially higher, due to the superposition of pile-up events.

The official draft trigger table for low luminosity is reproduced. This table allows capture of single jets with 134 GeV transverse energy at 95% efficiency. For missing E_t, 95% efficiency is achieved at 200 GeV.

Chapter 6 - Trigger Tower Size in the Very Forward Calorimeters

The lateral granularity of the forward calorimeter readout towers has been studied in detail, and an optimal tower size of 0.17×0.17 in (η, ϕ) has been chosen [60]. However, the granularity of the Level-1 trigger towers has yet to be decided. It has been suggested that the Level-1 towers should be no larger than the readout towers. Since the cost of trigger electronics rises with the number of trigger towers, though, the towers should be made as large as possible before the trigger efficiency for physics signals is compromised.

Two possible segmentation schemes have been investigated, where the Level-1 trigger towers correspond to 3×2, or 2×2, readout towers in (η, ϕ). The Level-1 response to jets has been simulated with both schemes in order to evaluate each scheme and quantify the differences between them. It is clear that these differences will be small, but a full simulation was felt to be necessary before making a decision.

6.1 Simulation

To reproduce the entire binned minimum-bias dataset for both trigger segmentations would require far more cpu time than was available. Instead, a small signal sample was used. The signal in question was that of an invisibly decaying Higgs boson produced via weak boson fusion, as used in Chapter 5 to measure the Level-1 missing E_t resolution and trigger turn-on. High luminosity (10^7 $mb^{-1}s^{-1}$) running conditions were simulated. Under these conditions, the average number of minimum-bias pile-up events per bunch crossing is 17.3.

Production of the data sample was a complex process. In ORCA, the reconstructed hits and trigger primitives are stored in the Objectivity database during ‘digitization’ (see§3.11.1). Since the trigger tower segmentation is implemented during digitization code, this stage must be run separately for each segmentation scheme. It would be relatively trivial to produce a separate database for each scheme, but this would require a large amount of disk space and cpu time that were not forthcoming. Instead, a *shallow copy* of the original database was made. A shallow copy is a database that simply contains pointers to data within the original database, rather than copies of the data itself. In this case, new trigger primitive data was generated for a different segmentation scheme, and the relevant pointers in the shallow copy were replaced with pointers to the new data. In this way, the different segmentation schemes could be accommodated without affecting users of the original dataset, and while minimizing the disk and cpu requirements. The original dataset was made at FNAL, but for various reasons, the shallow copy was made at

CERN. This is an example of GRID technology in use in High Energy Physics, where the user sees a database in one location, but does not know (or care) where the data is actually physically located.

Once the shallow copy had been filled with trigger primitives for each segmentation scheme, the Level-1 output was collected using the author's code described in §5.2.3. Ntuples were produced for both '3×2' and '2×2' segmentations and the histograms were produced using Fortran routines (as described in §5.2.3).

6.2 Jet Trigger Performance

The two schemes are evaluated by calculating the energy resolution, trigger turn-on and rate, for jet and missing E_t triggers. Since we are only interested in the Level-1 response from the Very Forward calorimeters, the forward jet stream is taken from the Level-1 output, and all other candidates are ignored. This is equivalent to demanding $3 < |\eta_{\text{jet}}| < 5$.

6.2.1 Jet Matching

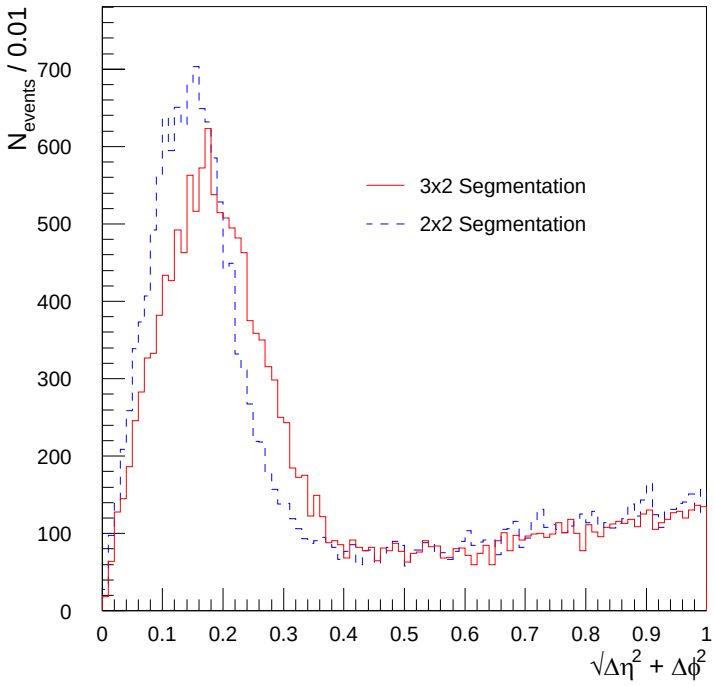

Figure 6.1 - Separation in (η, ϕ) of Level-1 forward jets and nearest generator jet.

In order to calculate energy resolution and trigger efficiency, we must find the generator level jet from which each Level-1 jet originated. This is done by finding the nearest generator jet in (η, ϕ) for each Level-1 jet, and applying a cut to the spatial separation. Figure 6.1 shows the

distribution of this quantity before the cut. As is to be expected, the position resolution of the '2 × 2' scheme is slightly better than that of the '3 × 2' scheme. The tail in both distributions at high separation is caused by fake Level-1 jets. The Level-1 jet algorithm will always find jets, from noise and pile-up if no real jets are present, but these jets are not correlated with real generator level jets, hence the tail in Figure 6.1. An E_t cut applied to Level-1 jets reduces the fraction of jets in the tail, but none was applied here as the choice of matching cut value is not affected. The matching cut used in this study for both schemes was $\sqrt{(\Delta\eta^2 + \Delta\phi^2)} < 0.4$.

6.2.2 Jet Energy Corrections

The jet energy is corrected, as described in §5.3.1 and reference [54], using a quadratic fit to the Level-1 jet E_t as a function of true jet E_t. As mentioned in §5.3.1, a linear fit described the forward jet data well for low luminosity, but the effects of pile-up at high luminosity mean a small quadratic term improves the fit (as is the case for central jets at low luminosity). Details of the fit and the resulting correction coefficients are given in Appendix B.

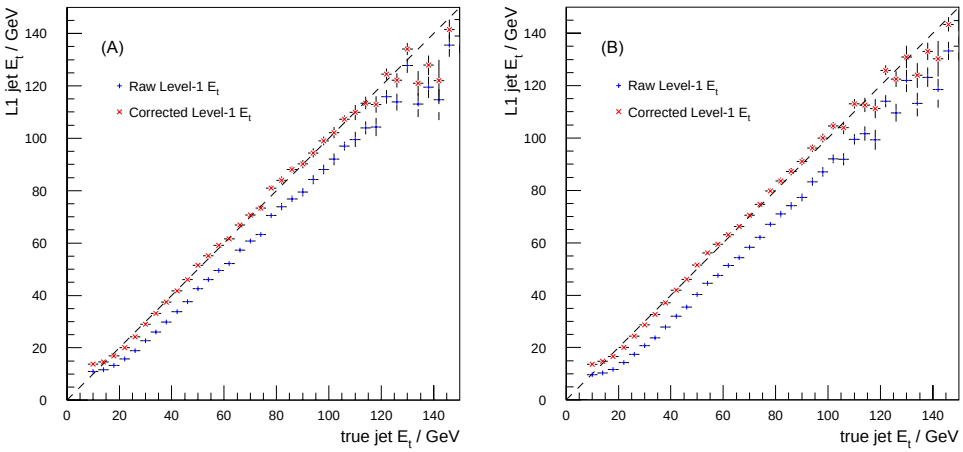

Figure 6.2 - Level-1 jet E_t as a function of true E_t before and after corrections, for (A) '3 × 2' and (B) '2 × 2' segmentation schemes.

The corrected and uncorrected Level-1 jet E_t is plotted as a function of true jet E_t for the '3 × 2' segmentation scheme in Figure 6.2a and for the '2 × 2' scheme in Figure 6.2b. The equation $E_t^{L1} = E_t^{true}$ is shown on each figure with a dotted line. As can be seen the corrected energy follows the line of equality well, up to energies of ~125 GeV. Deterioration of the corrected E_t beyond this point is to be expected, since the sample contains insufficient jets with E_t higher than ~150 GeV.

6.2.3 Transverse Energy Resolution

Figure 6.3 shows the E_t resolution for matched Level-1 forward jets as a function of true jet

E_t, for each segmentation scheme. The resolution is calculated after the energy corrections described above have been applied. A standard parameterization of hadronic calorimeter resolution (see Eq. 6.1) is fitted to the data using a least squares method.

$$\frac{\sigma_{Et}}{E_t} = \frac{p_1}{\sqrt{E_t}} \oplus p_2^2 \oplus \frac{p_3}{E} \tag{6.1}$$

The fitted curves are shown in Figure 6.3, and the resulting parameters are listed in Table 6.1. These figures indicate an improvement in the fractional resolution of around 2% going from '3×2' to '2×2' segmentation (i.e. for jets with 100 GeV E_t, the resolution goes from ~19% to ~17%). However, as can be seen from Figure 6.3, the improvement is really obtained for relatively low E_t jets, and above 50 GeV there is little difference between the two schemes.

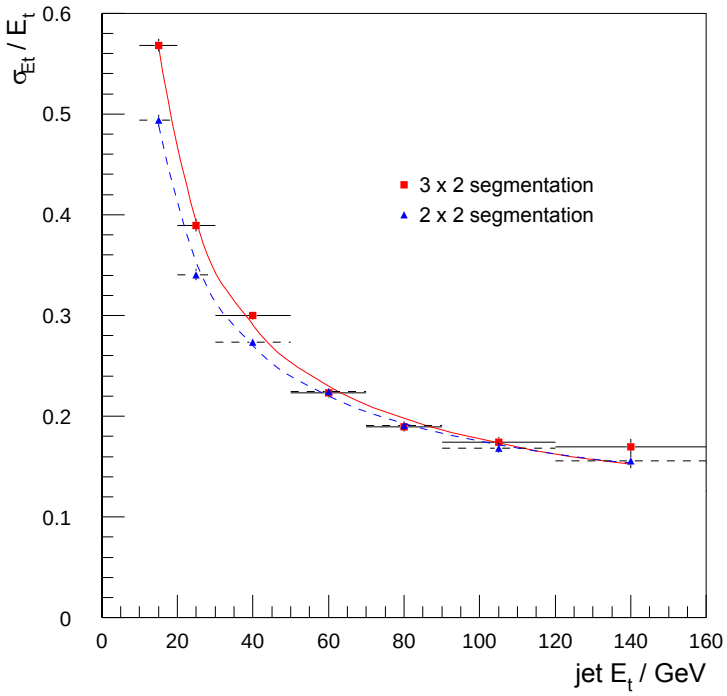

Figure 6.3 - Jet E_t resolution, including fitted curves, for '3×2' and '2×2' schemes (solid and dashed lines, respectively).

Segmentation	p_1 / GeV$^{1/2}$	p_2	p_3 / GeV
'3×2'	1.49 ± 0.25	0.07 ± 0.06	6.1 ± 0.7
'2×2'	1.46 ± 0.21	0.09 ± 0.04	4.5 ± 0.8

Table 6.1 - Resolution fit results. Removing the p2 parameter resulted in a degradation of chi-squared.

This improvement for low E_t jets can be explained in terms of pile-up, and the solid angle subtended by the jet. The larger '3×2' scheme samples more pile-up than the '2×2' segmentation, and hence is more susceptible to fluctuations. This manifests itself as an improvement in the resolution. However, since the energy scale of the pile-up is relatively low, as a fraction of total energy measured the fluctuations are only significant for low energy jets. This is corroborated by the fit results. The stochastic and constant terms remain the same for both segmentation schemes, but the noise term is reduced in the '2×2' scheme. The noise term is theoretically proportional to the amount of noise sampled by the jet algorithm. Indeed, the ratio of noise terms for the two schemes here is roughly equal to the ratio of Level-1 jet areas in (η, ϕ).

6.2.4 Trigger Turn-on

The resolution results obtained above would suggest slightly better turnon for jets below 50 GeV in the 2×2 scheme, and little difference for jets above this threshold. Turn-on curves were calculated using the corrected jet energy, and are shown in Figure 6.4 for '3×2' and '2×2' segmentations. Little difference, if any, can be inferred between the two sets of curves as they sit almost on top of one another. If one looks closely at the first 3 plots in Figure 6.4, one can almost detect a sharpening of the turn-on in the '2×2' case. The resolution results indicate sharpening in the turn-on curve to the order of a few per cent, at most 10%, or in the case of the 30 GeV L1 threshold, say, no more than 3 GeV. This is consistent with what can be observed in Figure 6.4.

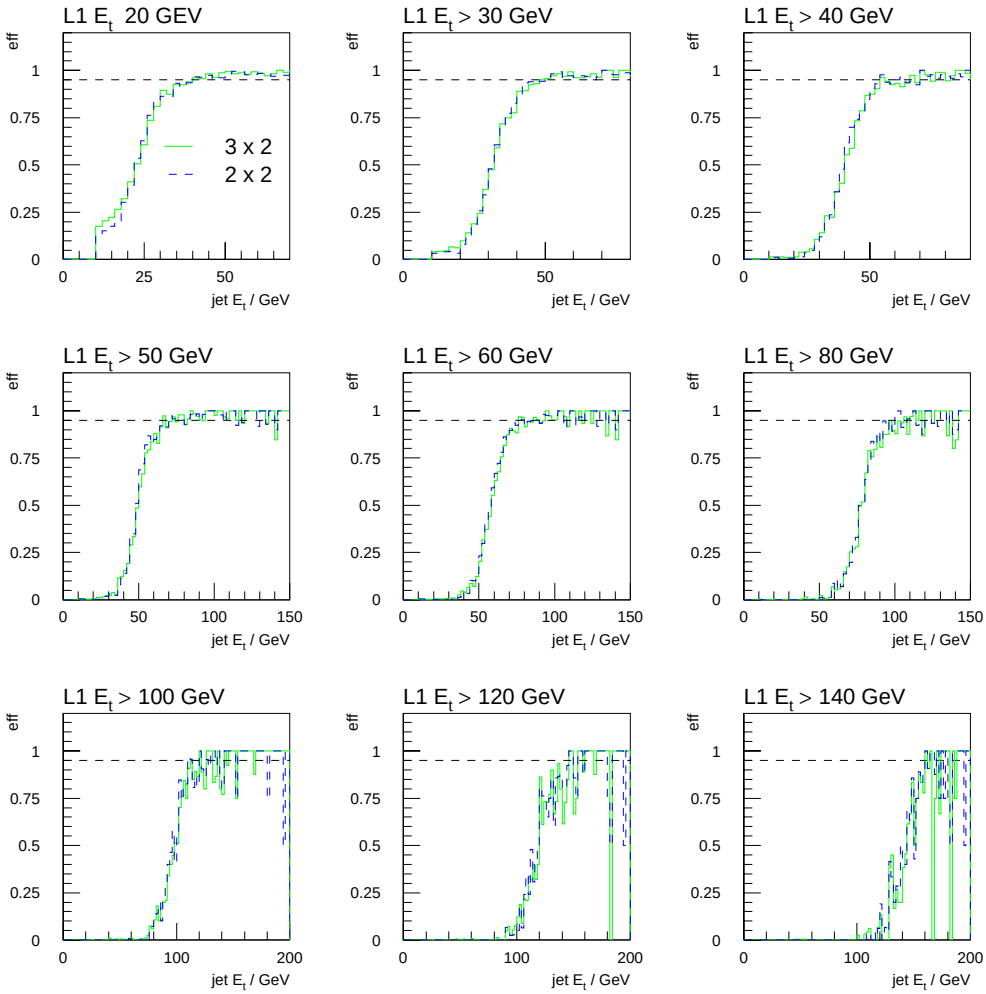

Figure 6.4 - Comparison of turn-on curves for both segmentation for lower Et thresholds. The 3 x 2 segmentation is shown with a solid line, and the 2 x 2 segmentation is shown with the dashed line.

6.2.5 Trigger Rate

Clearly, the true jet rate cannot be calculated for the '2×2' segmentation because the binned minimum-bias samples are not available, but a comparison between the two segmentation schemes can be made.

Figure 6.5 shows the number of events in the sample that pass single jet, di-jet, tri-jet and quad-jet triggers as a function of E_t threshold. The single and di-jet triggers show a slightly higher rate for the '2×2' segmentation than the '3×2' segmentation, for E_t thresholds above ~60 GeV. The tri- and quad-jet triggers, however, show a more significant difference, with the '3×2' segmentation giving the higher rate in both cases.

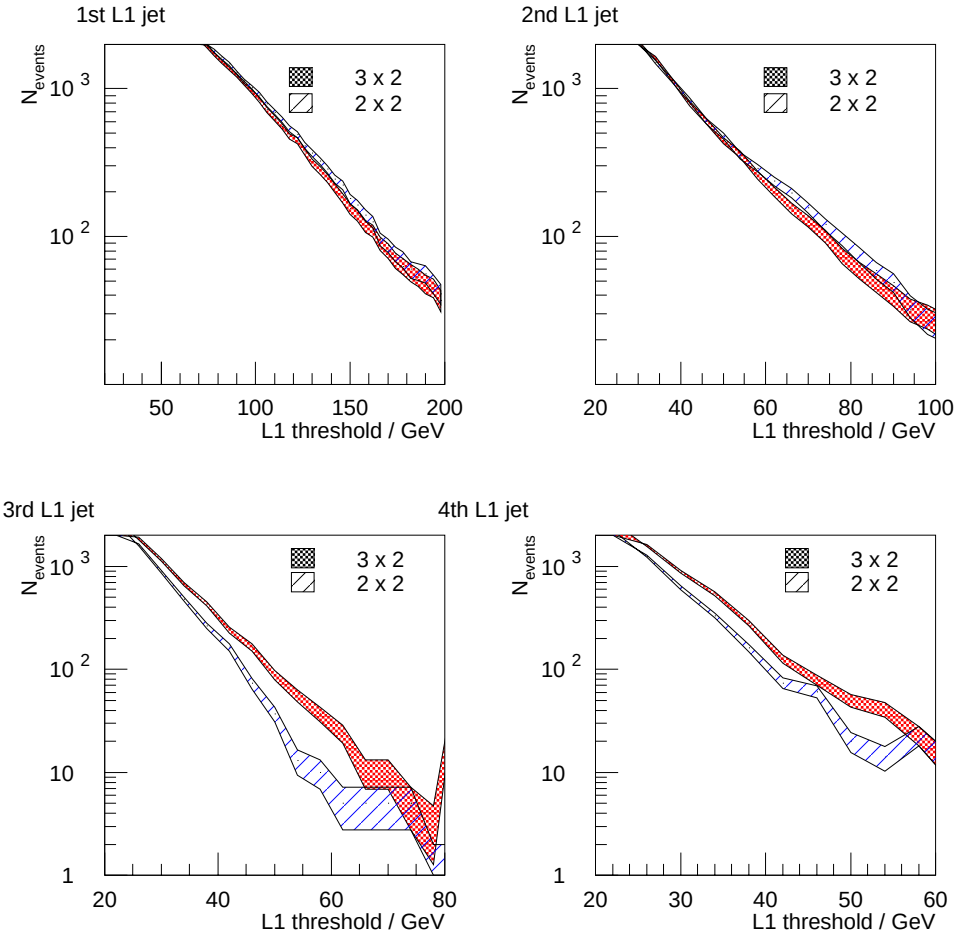

Figure 6.5 - Jet trigger 'rate' for single, di-, tri- and quad-jet triggers.

The difference in single jet trigger rate between the two segmentation schemes is apparent well above the calibration region. It is likely, therefore, that the difference is caused by miscalibration. Although the di-jet rate difference is clear above 60 GeV, which is inside the good calibration region, the separation between the two schemes is very slight, and may also be due to miscalibration.

The difference in tri- and quad-jet trigger rate between the two schemes is, however, more substantial. The rate is affected by the number of trigger towers, the energy resolution, and the number of 'fake' jets. These three effects are discussed in the next three paragraphs.

Due to its simple jet definition of one trigger tower having greater energy than its neighbours, the Level-1 jet algorithm, in the presence of noise and pile-up particles, will always find jets. Since the finer-grained scheme has a larger number of possible jets, due the increased number

of trigger towers, it will find more jets than the coarser scheme.

The improved resolution of the finer-grained segmentation at low energies will reduce the rate slightly. As discussed in §5.5.2, the Level-1 rate results from the convolution of the 'true' rate with the Level-1 resolution, so a tighter resolution results in a lower rate. Again we may estimate the magnitude of this effect. The jet E_t resolution for 40 GeV jets with the '3×2' segmentation is 12.0 GeV, and with the '2×2' segmentation is 10.8 GeV. So the difference in rate due to this effect alone is roughly equivalent to a change in Level-1 threshold of ~1 GeV. In other words, the trigger rate for the '3×2' segmentation with a threshold of 40 GeV is approximately the same as the rate for the '2×2' segmentation with a threshold of 41 GeV.

It is suggested that the tri- and quad-jet rate difference between the two schemes is due to improved rejection of 'fake' jets. A fake Level-1 jet may arise either where there is no real jet at all, or where several soft jets from pile-up events deposit energy in the same calorimeter region. The former case will generally result in reasonably low energy fake jets, but the latter can result in fake jets with high energy, despite the fact that the contributing jets are all soft. The '2×2' scheme probably finds the same number of fake jets, as the '3×2' scheme, but because the solid angle over which the jet energy is calculated is smaller, it should result in a fake jet with lower energy.

The rate difference for the tri- and quad-jet triggers appears to be approximately the same as a 5-10 GeV change in Level-1 threshold. While this is a reasonably large effect at low energies, at reasonable thresholds it is less significant.

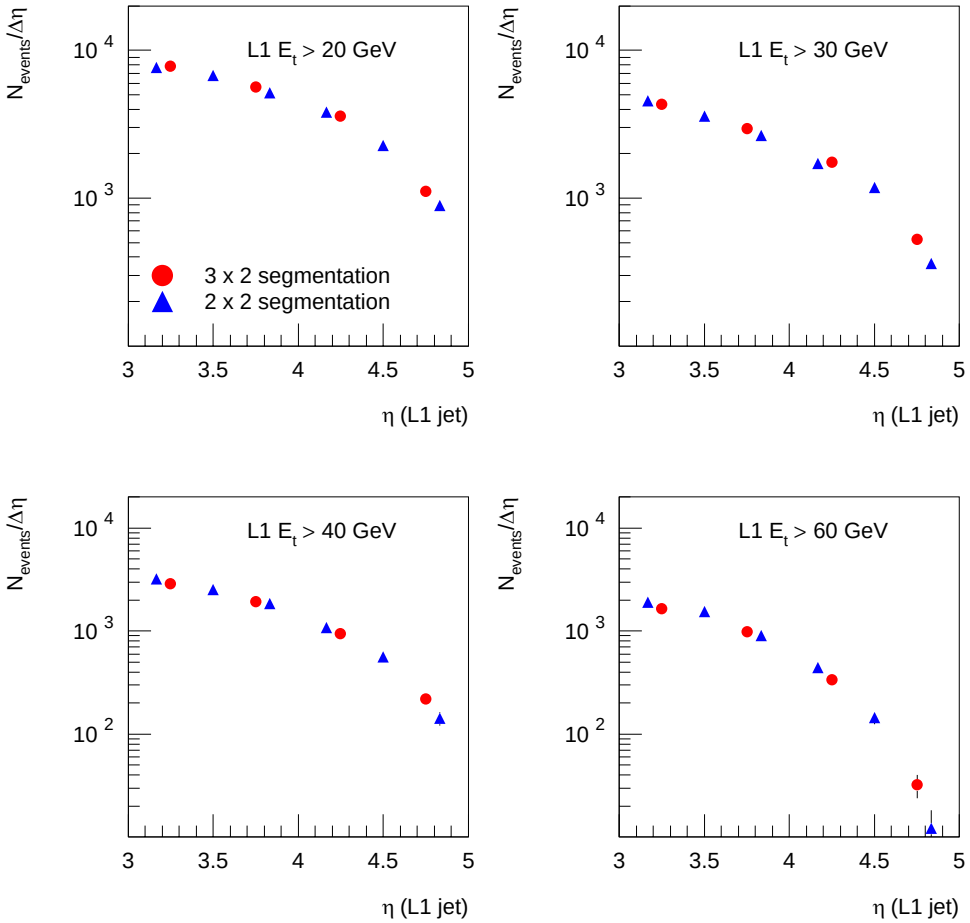

Figure 6.6 - Jet trigger 'rate' per unit pseudorapidity plotted as a function of $|\eta|$ for several E_t thresholds in both segmentation schemes.

Figure 6.6 shows the jet rate as a function of pseudorapidity for four different E_t thresholds. The rate for each segmentation is scaled according to the tower width in pseudorapidity to account for the true flux of particles as a function of pseudorapidity. No significant difference can be observed between the two segmentation schemes from these plots.

6.3 Missing E_t Trigger Performance

The Level-1 missing energy calculation is performed by summing the X and Y components of transverse energy over trigger towers. Since the towers are of non-zero size, the transverse energy is effectively averaged over the tower during the calculation. So the missing energy calculation is affected by the tower size. The effect is expected to be very slight since the change in tower size is small, but the missing E_t resolution was calculated nonetheless.

The fractional E_t^{miss} resolution is plotted as a function of true E_t^{miss} in Figure 6.7. Again, Eq.

6.1 is fitted to the data using a least squares method. The resulting fit parameters are given in Table 6.2. The resulting fit parameters show a very slight improvement going from '3×2' to '2×2' segmentation. Again, the improvement is essentially only obtained at low values of missing E_t. Above 60 GeV, the two schemes cannot be distinguished in these terms.

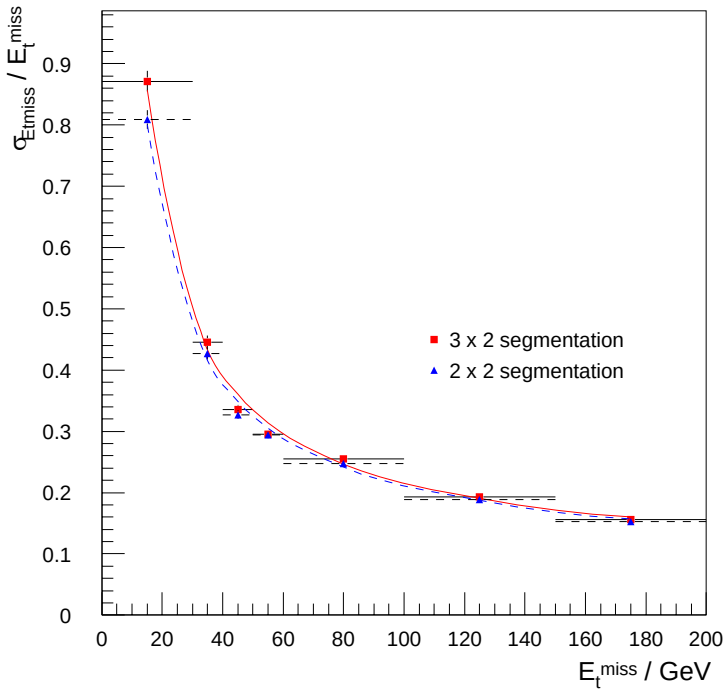

Figure 6.7 - E_t^{miss} resolution for '3×2' and '2×2' tower schemes

Segmentation	p_1 / GeV$^{1/2}$	p_2	p_3 / GeV
'3×2'	1.73 ± 0.10	0.07 ± 0.02	10.9 ± 0.4
'2×2'	1.79 ± 0.09	0.06 ± 0.02	9.7 ± 0.4

Table 6.2 - E_t^{miss} resolution fit parameters

6.4 Summary & Conclusion

Clearly, there are some slight gains to be found in using the '2×2' segmentation over the '3×2' segmentation. The most significant effect is an improvement in the E_t resolution of low E_t jets. If the towers under consideration were the readout towers, that dictate the final resolution of offline jets, perhaps this improvement would justify the added cost. At Level-1, though, the foremost criteria is efficiency, and little or no improvement is observed in this area. The finer granularity scheme may offer improved rejection of fake jets, but again, only at low energies.

Finally, a very slight improvement in the missing E_t resolution, at low values of missing E_t, is observed going from the '3×2' to '2×2' scheme.

To conclude, the '3×2' scheme is found to offer no significant degradation in performance over the '2×2' scheme, and due to its lower cost, is the preferred solution.

Chapter 7 - Topological Di-jet Triggers

This chapter describes a study of the use of tag jets from weak boson fusion (WBF) processes in the Level-1 trigger. WBF processes include a pair of hard forward jets in the final state (see Figure 7.3), and the purpose of this study was to investigate how far the di-jet trigger can be extended by applying topological cuts to the jet pair. These triggers are evaluated both for WBF processes in general, and for a Higgs produced in WBF decaying to invisible particles. The most appropriate triggers are selected in both cases, and the total Level-1 trigger efficiency is estimated.

7.1 Introduction

For some signals, it will be sufficient to use a Level-1 trigger consisting simply of a threshold on the highest jet, or missing E_t measurement. In other channels, however, the thresholds that are acceptable to the DAQ system in terms of rate may provide inadequate signal efficiency for physics analysis. In such cases, more sophisticated triggers must be constructed. One possible enhancement, available in the Global Trigger logic, is the application of topological cuts.

The physics signals under consideration here are those of a Higgs produced in weak boson fusion. These signals produce two ‘tag’ jets in the forward direction, providing a useful signature for rejection of backgrounds in the offline analysis. The Level-1 trigger signature had previously relied purely on the decay products of the Higgs, which presents no problem for a massive Higgs decaying to four leptons, for example, but proves problematic in other cases; e.g. a light Higgs decaying invisibly. If it is possible to include the tag jets in the Level-1 trigger, it is expected that the trigger efficiency of such signals can be improved.

The primary goal of this study was to produce a trigger based solely on tag jets, that could be used to trigger on WBF processes in general, regardless of the Higgs decay mode. This would be of use both in cases where the decay products present a more difficult challenge for the trigger, and as a redundant trigger for other decay modes. In the latter case, the trigger efficiency can be estimated with greater accuracy. However, in the event that such a trigger is unfeasible, the tag jets may still be of use in conjunction with the Higgs decay products, in a trigger that will provide improved performance over the decay products alone.

This chapter discusses the identification of the tag jets at Level-1, and how they may be used to increase trigger performance for WBF signals. First, the kinematics are studied at generator level. Here, cuts on the separation of the tag jets in pseudorapidity and azimuthal angle are

identified which may enhance the efficiency of the di-jet trigger. The reconstruction of the tag jets at Level-1 is then shown, together with the reproduction of the separation distributions. The initial aim was to produce a trigger that uses tag jets alone to capture WBF signals. The performance of such a trigger is found to provide insufficient efficiency at reasonable rates. The tag-jet + missing energy trigger was investigated for WBF processes where the Higgs decays invisibly. This production mode provides the best chance of observing invisible Higgs decays at the LHC [61]. The event selection efficiency for a 'di-jet + missing energy' trigger with a cut on the jet separation in pseudorapidity is found to provide excellent efficiency at a reasonably low rate.

The simulation techniques and monte-carlo samples used in this study were all described in Chapter 5. The invisible Higgs sample was used to measure the efficiency of the triggers, while the binned minimum-bias samples (with $\hat{p}_t$ in the range 10 - 800 GeV) were used to calculate the trigger rate. Quantities associated with the WBF tag quarks, and the Higgs itself, were calculated using information from the PYTHIA output, stored and retrieved with ORCA.

7.2 Weak boson fusion

The cross-section of various Standard Model Higgs production modes at the LHC is shown in Figure 7.1 as a function of Higgs mass. The Feynman diagrams for two of these modes, gluon-gluon fusion and weak boson fusion, are shown in Figure 7.2 and Figure 7.3, respectively. Gluon-gluon fusion is the dominant production mode, but as can be seen from the Feynman diagram, the only final state particles in this case are the decay products of the Higgs. Weak boson fusion, however, is a production process that does provide a signature. Even if the Higgs decay products are invisible to the detector, this signal may be isolated using the 'tag' jets associated with the final state quarks.

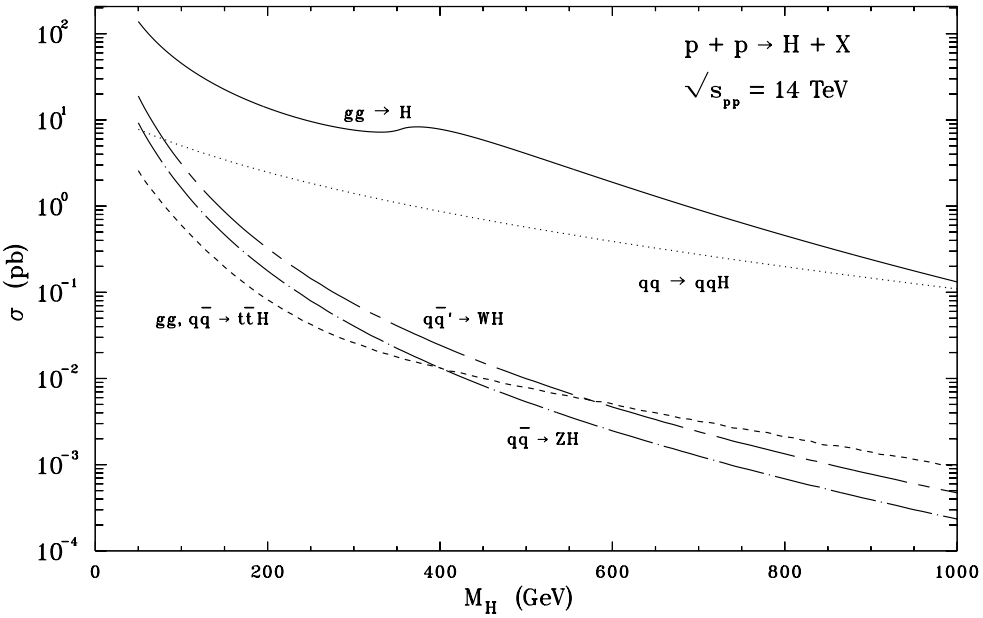

Figure 7.1 - Cross-section of various Higgs production processes, as a function of Higgs mass (taken from [62] - produced well before the current upper limits on m_H were calculated)

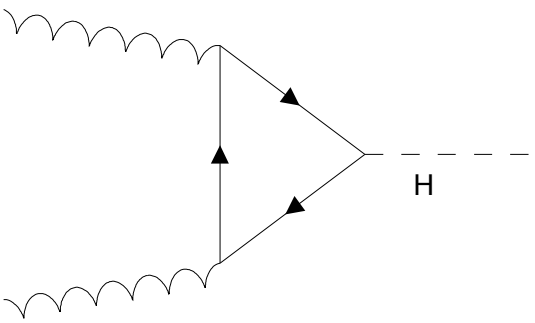

Figure 7.2 - Feynman diagram of a Higgs produced via gluon-gluon fusion

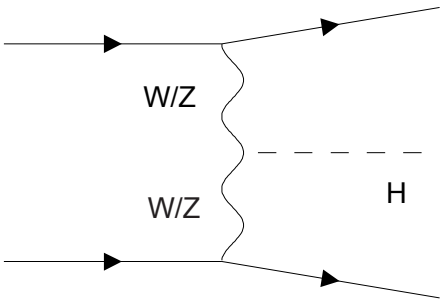

Figure 7.3 - Feynman diagram of a Higgs produced via weak boson fusion.

Various Higgs discovery channels using WBF have been suggested. Most of the analyses for gluon-gluon fusion production can be adapted to WBF, although the advantage gained from using the tag jets does not always outweigh the lower cross-section. The most noteworthy WBF channels are those involving the following Higgs decays:

- $H \rightarrow W^+ W^- \rightarrow l^+ \nu \;\; l^- \overline{\nu}$
- $H_{SUSY} \rightarrow \tau^+ \; \tau^-$
- $H \rightarrow$ invisible particles (e.g. $\chi_1^0 \chi_1^0$ in the MSSM)

7.2.1 Signal Characteristics

The transverse energy distributions of the two tag quarks and the Higgs in WBF are shown in Figure 7.4. As can be seen, the tag quarks are reasonably hard. However, it should be clear that only a small proportion of the sample will pass the jet triggers given in Table 5.4.

The distribution of the quarks and Higgs in pseudorapidity are shown in Figure 7.5. It is worth noting that the tag quark distribution is more forward than that of QCD di-jets (see Figure 5.2 for comparison).

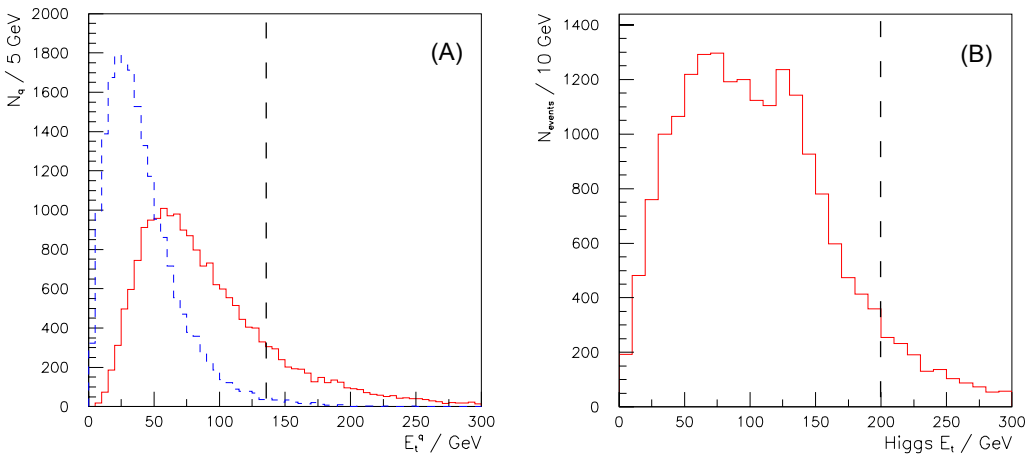

Figure 7.4 - Transverse energy of (a) the tag quarks and (b) the Higgs in WBF. The Level-1 95% efficiency points for 110 GeV single jet and 105 GeV E_t^{miss} triggers are marked with dashed lines.

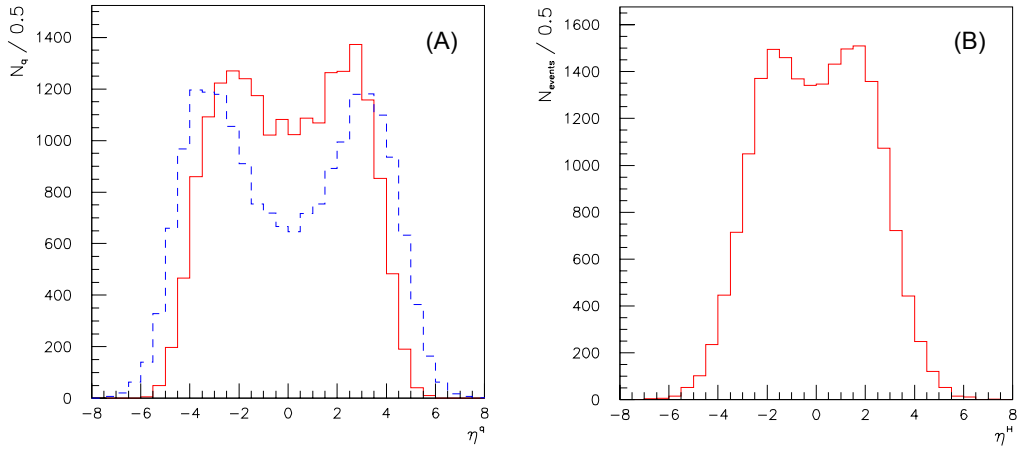

Figure 7.5 - Distributions in pseudorapidity of (a) the tag quarks and (b) the Higgs in WBF

It is useful at this point to define two simple cuts (that will be referred to later as the 'tag quark' cuts) that select the useful event sample. We cannot reconstruct jets outside these limits with reasonable efficiency, so event failing these cuts are of little use for analysis.

$$|\eta| < 5 \tag{7.1}$$

$$E_t > 20GeV \tag{7.2}$$

These cuts essentially define the limits of what can be reconstructed using the full offline algorithms, and are passed by 75.6% of the WBF sample. It should be noted, though, that individual WBF analyses may use a higher E_t cut on the offline jets in order to improve the rejection of physics backgrounds.

7.2.2 Baseline Jet Trigger Performance

The performance of Level-1 single jet and di-jet triggers in selecting WBF events is shown in Figure 7.6. Here, the abscissa is given by WBF selection efficiency and the ordinate is given by trigger rate, for various values of threshold. The efficiency is calculated with respect to the 'tag quark' cuts, described above. It is clear that even if a large fraction of the DAQ bandwidth is available for the single jet trigger the efficiency that can be achieved is low (merely 50% at 10 kHz).

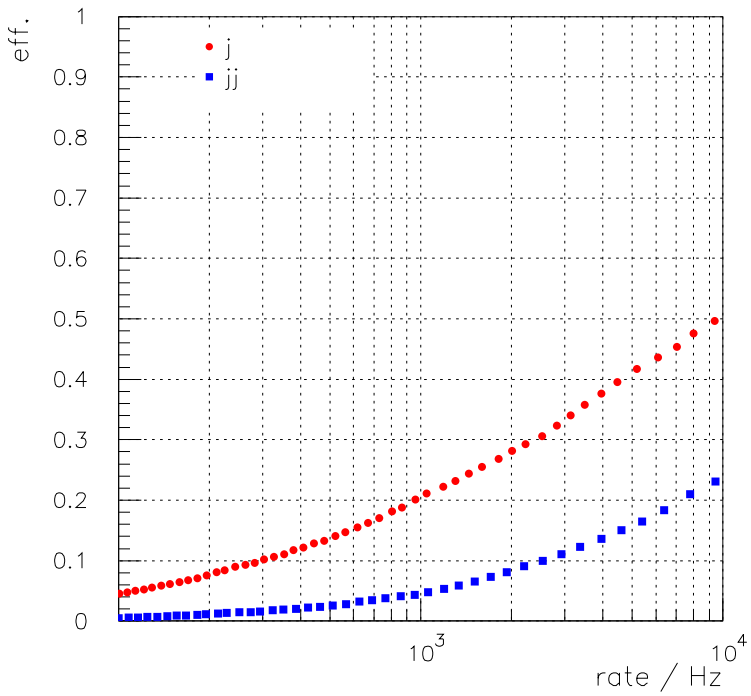

Figure 7.6 - Level-1 trigger rate vs. WBF selection efficiency for single jet and dijet triggers. Efficiency is calculated with respect to the tag quark cuts.

7.2.3 Tag Quark Correlation

As is to be expected, the relationship between the two tag quarks is very different from that of QCD dijet production. Figure 7.7 shows the spatial separation in η and ϕ for WBF and QCD di-jets. The separation is calculated between quarks for the WBF sample, and between the two highest E_t generator level jets for the QCD sample.

The $\Delta\eta$ plot clearly shows the increased separation of the WBF quarks over the QCD di-jets, while the $\Delta\phi$ plot shows the QCD di-jets are produced back-to-back, as required by conservation of momentum in the transverse plane. Since the WBF process has 3 bodies in the final state, the $\Delta\phi$ distribution of the tag quarks does not exhibit this behaviour. Incorporation of suitable topological cuts with the Level-1 di-jet trigger, requiring large $\Delta\eta$ and/or $\Delta\phi$ significantly smaller than π, may allow the E_t threshold to be lowered sufficiently to capture a useful proportion of the sample.

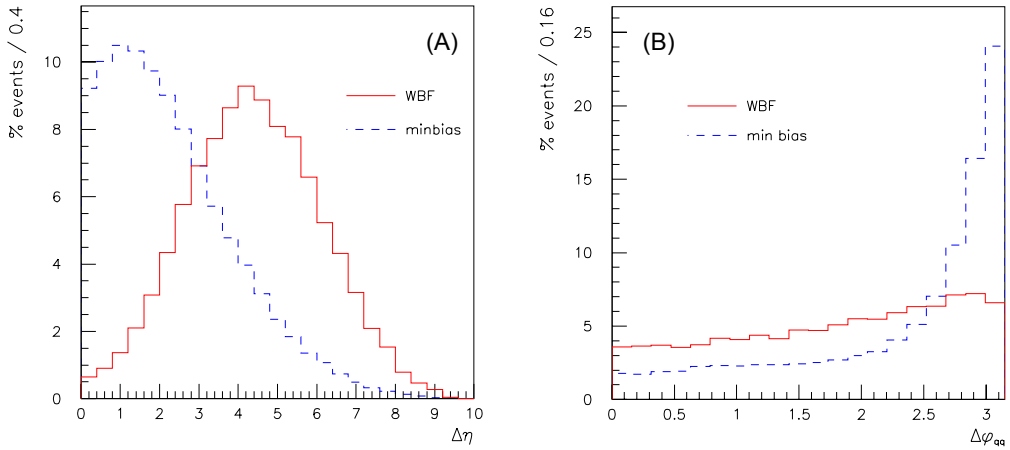

Figure 7.7 - Spatial separation of tag quarks in (A) pseudorapidity and (B) ϕ-angle for WBF signals. Separation of top two jets in minimum bias events is also shown for comparison.

In order to find the optimal cut value, we plot the ratio $S/\sqrt{B}$ where S and B are respectively the number of signal and background events passing the generator level cut. Figure 7.8 shows this ratio, for 10 fb^{-1} integrated luminosity, as a function of cut value for $\Delta\eta$ and $\Delta\phi$ cuts. The actual values given for $S/\sqrt{B}$ are very low because no threshold has been applied to the energy of the quarks/jets. It is the shape of the distributions that we are more interested in here. As can be seen, the maximal values for $S/\sqrt{B}$ are given by cuts of $\Delta\eta > 4$ and $\Delta\phi < 2.5$.

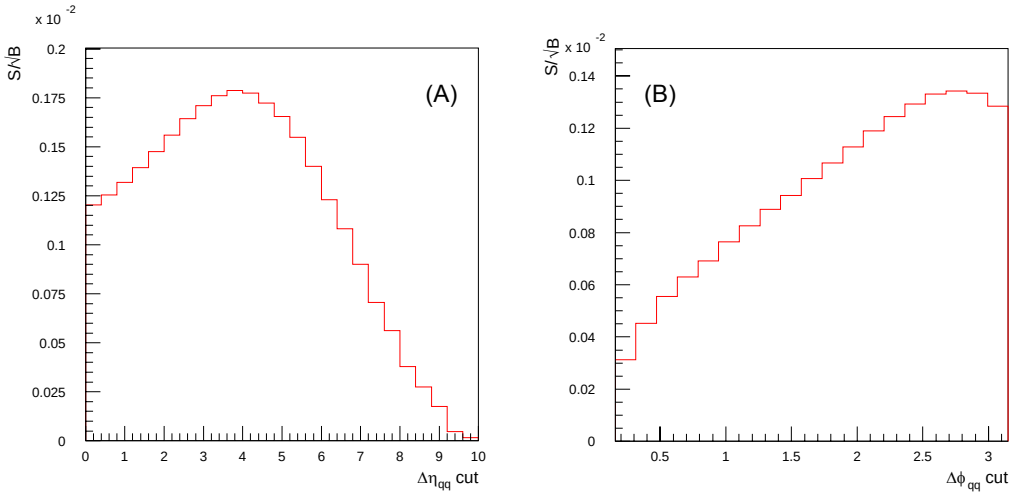

Figure 7.8 - Signal (WBF) to background (min. bias) ratio for (A) $\Delta\eta$ and (B) $\Delta\phi$ cuts on tag quarks. No E_t threshold is applied.

7.3 Reconstruction in the Level-1 Trigger

In order to make full use of the spatial correlation between the tag quarks, we must reproduce this information in the trigger with good efficiency. The extent to which the Level-1 trigger achieves this is evaluated in this section.

7.3.1 Tag Jet Identification

For the purposes of this study, the separate streams of central, tau and forward jet candidates are ignored, and the 12 jets available in the GT are simply sorted by E_t. Although this single jet list does not represent the exact logic planned for the hardware, this work was carried out during a period when the exact nature of the candidate streams available to the GT was in a state of flux. The same effect can be achieved using the 3 separate jet lists, by implementing the triggers listed in §7.3.3 using the 6 necessary jet pairs and taking the logical OR of the result. Clearly this uses up more logic in the GT than if the jets are available in a single stream, but it is nevertheless possible. The efficiency of such a scheme will remain the same as that presented here, but theoretically the trigger rate may increase. Using the results shown in §7.4.4, it will be argued that any such increase will be insignificant. Finally, if little or nothing can be achieved with a single jet list, then the '3 separate list' scenario represents no improvement.

The result of matching the two highest E_t Level-1 jets to the tag quarks is shown in Figure 7.9A. Here, the solid line shows the position of the highest E_t quark and the dashed line shows the position of the other quark. The matching procedure used was as described in §5.2.5, using $\Delta R < 1$ as the matching criterion. As can be seen, the first tag quark is selected by the first L1 jet in more than 80% of WBF events, and the 2nd tag quark is selected by the 2nd L1 jet in nearly 70% of WBF events.

The efficiency with which a particular pair of Level-1 jets (or the single highest E_t jet) will select the two tag quarks is given in Table 7.1. As is to be expected, the single highest E_t jet selects a tag quark with greater efficiency than the any pair of Level-1 jets select both tag quarks. It is also clear that using more than the first three jets to identify the tag quarks will provide little gain.

Figure 7.9B shows the efficiency for the 1st Level-1 jet matching the 1st tag quark as a function of the tag quark E_t, together with the efficiency for 1st & 2nd Level-1 jets matching both tag quarks as a function of the lower tag quark E_t. Naturally, the harder quarks are easier

to identify, and the situation will improve as the Level-1 threshold increases.

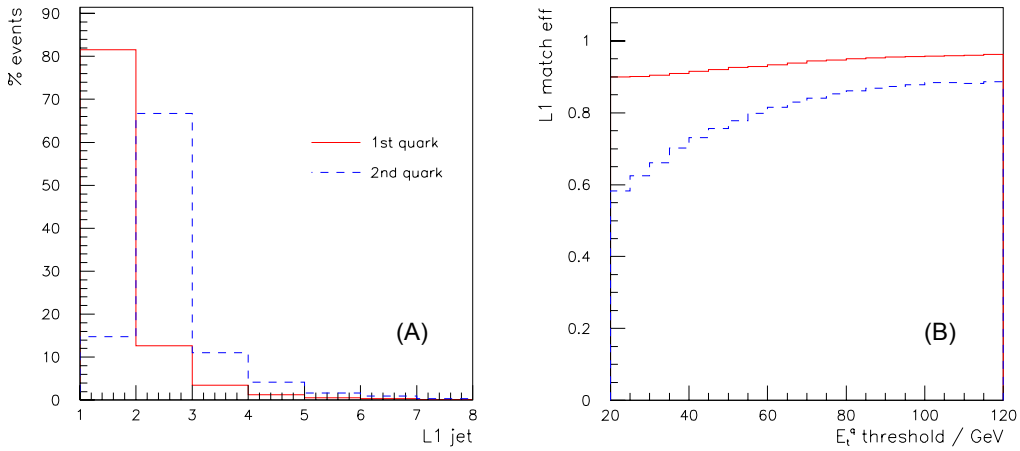

Figure 7.9 - Level-1 tag quark selection. (A) shows where in the L1 output the tag jets appear. (B) shows efficiency for finding tag quarks in 1st position (solid line), and in both 1st & 2nd position (dashed line).

Level-1 jet(s)	Eff. ($\Delta R < 0.5$)	Eff. ($\Delta R < 1.0$)
1	90%	93%
1 & 2	58%	67%
1 & 3	5%	7%
1 & 4	1%	2%
2 & 3	1%	1%

Table 7.1 - Tag quark selection efficiency by Level-1 jet trigger, for different matching cuts.

7.3.2 Topological Cuts at Level-1

Figure 7.10 shows the spatial separation of the top two Level-1 jets for minimum bias events and for a WBF signal, after applying a cut on both jets of $E_t > 30$ GeV. Comparison with Figure 7.8A shows the Level-1 tag quark selection efficiency is sufficient to reproduce the $\Delta\eta$ distributions reasonably well. The same cannot be said for the $\Delta\phi$ distributions, however.

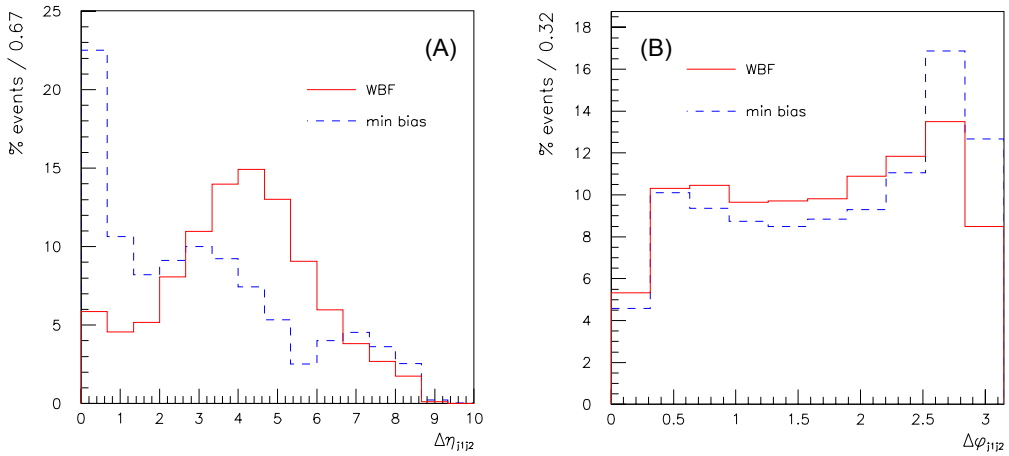

Figure 7.10 - Spacial separation of 1st and 2nd Level-1 jets, in eta and phi, for WBF (solid line) and minimum bias events (dashed line) after a cut on the L1 E_t of 30 GeV.

The value of $S/\sqrt{B}$ for 10 fb^{-1} integrated luminosity is shown for cuts on $\Delta\eta_{jj}$ and $\Delta\phi_{jj}$ at Level-1 in Figure 7.11. The values given here are rather better than for the generator level plots, as a nominal E_t threshold of 30 GeV has been applied. From this plot, we take the initial cut values to be:

$$\Delta\eta_{jj} > 3.5 \tag{7.3}$$

$$\Delta\phi_{jj} < 2.5 \tag{7.4}$$

This should provide optimal background rejection, while still retaining good signal efficiency.

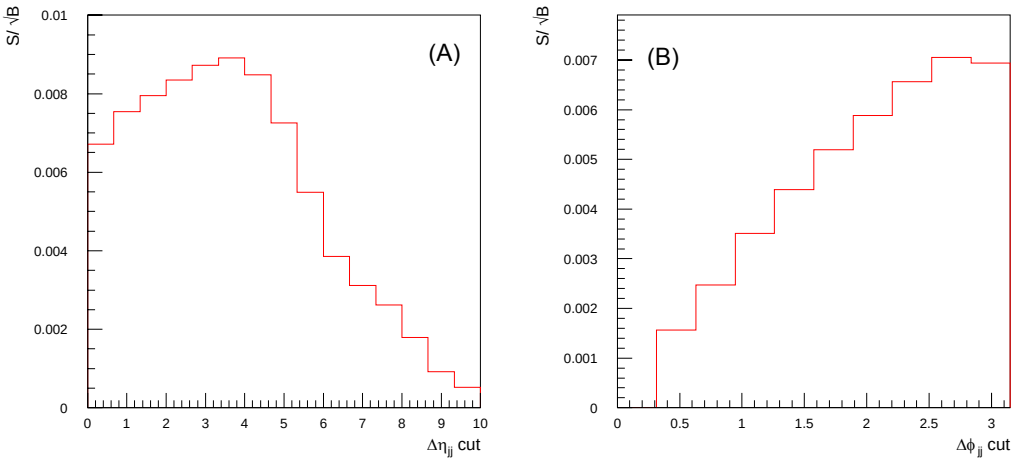

Figure 7.11 - Signal (WBF) to background (min. bias) ratio for $\Delta\eta$ and $\Delta\phi$ cuts at L1, after a 30 GeV di-jet E_t threshold. Again, the vertical scale is in arbitrary units.

Figure 7.12 shows the tag quark separation after application at Level-1 of an E_t threshold of 30 GeV and topological cuts. Figure 7.12A shows $\Delta\eta_{qq}$ after a Level-1 cut of $\Delta\eta_{jj} > 3.5$, while Figure 7.12B shows $\Delta\phi_{qq}$ after a Level-1 cut of $\Delta\phi_{jj} < 2.5$. As can be seen, the $\Delta\eta$ cut provides a reasonably sharp cutoff on the generator level quantity. However, the effect of the $\Delta\phi$ cut on the generator level quantity is rather less pronounced. This plot suggests that the Level-1 quantity is poorly correlated with the true value, and that the $\Delta\phi$ cut will not therefore provide background rejection as good as the $\Delta\eta$ cut.

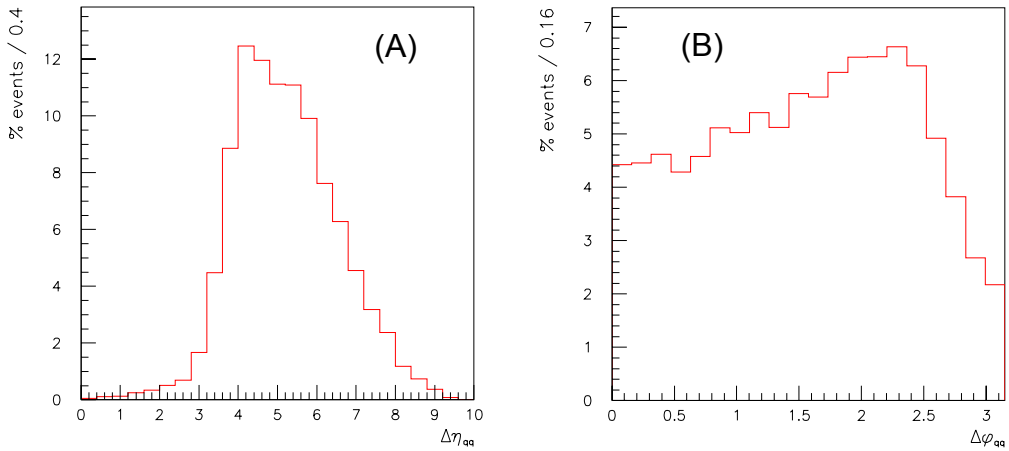

Figure 7.12 - A) Tag quark $\Delta\eta$ and B) $\Delta\phi$, after corresponding cuts at Level-1 ($\Delta\eta > 3.5$ in plot A, $\Delta\phi < 2.5$ in plot B.

7.3.3 Global Trigger Algorithms

The Global Trigger conditions considered here are:

$$E_t(J_2) > X + \Delta\eta(J_1J_2) > 3.5 \tag{7.5}$$

$$E_t(J_2) > X + \Delta\phi(J_1J_2) < 2.5 \tag{7.6}$$

$$E_t(J_2) > X + \Delta\eta(J_1J_2) > 3.5 + \Delta\phi(J_1J_2) < 2.5 \tag{7.7}$$

where the E_t threshold, X, remains a tunable parameter that can be used to control the rate. The values of the $\Delta\eta$ and $\Delta\phi$ cuts can be tuned at a later stage if the triggers are found to be useful. The values given above are felt to offer reasonable indication of whether the triggers are likely to perform to the desired standard.

7.4 Performance of Topological Triggers

As was shown in Figure 7.6, the standard di-jet trigger provides insufficient rejection of minimum-bias events to provide reasonable efficiency for WBF signals at an acceptable rate. However, as shown in the previous section, the selection of WBF events above the dijet background may be improved by including cuts on $\Delta\eta_{jj}$ and $\Delta\phi_{jj}$. In this section, the performance of pure tag jet triggers are described, in terms of rate and signal selection efficiency. The goal is a trigger based solely on tag jets that can be used to capture WBF signals. However, the results presented here suggest that this is unlikely to be feasible.

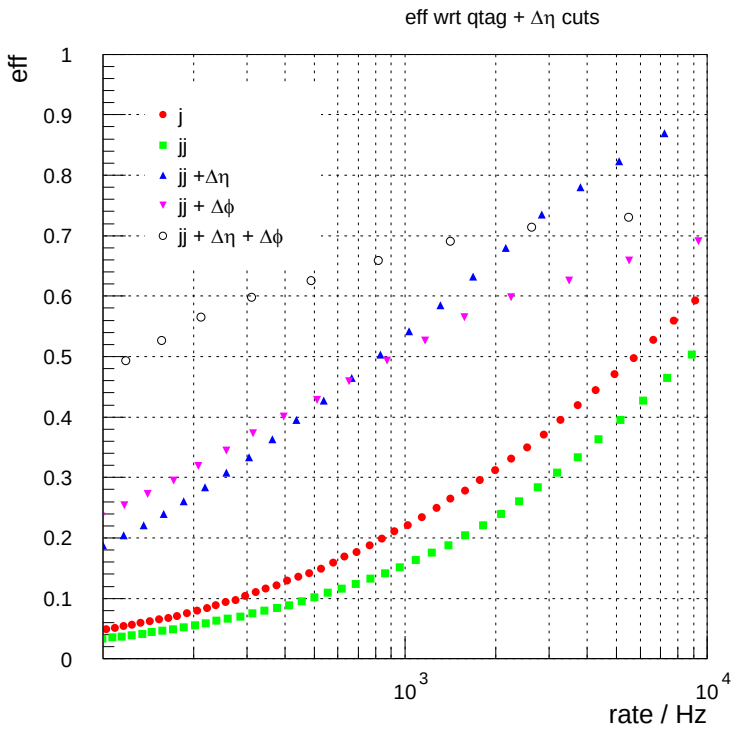

Figure 7.13 - Rate vs. efficiency for generator level 'triggers'. Efficiency is w.r.t the WBF sample after 'tag quark' and '$\Delta\eta$' cuts.

7.4.1 Generator Level Study

In order to estimate the potential performance of triggers including topological cuts, Figure 7.13 shows trigger rate vs. efficiency at the generator level. Here, the efficiency is calculated with respect to the tag quark and $\Delta\eta$ cuts, so this figure can be directly compared with Figure 7.16. It is clear that for a given trigger rate, the addition of topological cuts to the di-jet trigger provides a significant increase in efficiency over the single jet trigger. In the region of greatest interest, below 1 kHz, the combination of both cuts provides the best efficiency, by a significant margin, but either of the single cuts offers an improvement over the single jet

trigger. It should be noted that both the $\Delta\eta$ and $\Delta\phi$ triggers provide similar performance in this region, despite the fact that the efficiency is shown here with respect to a cut on $\Delta\eta$.

7.4.2 Trigger Rate

Figure 7.14 shows the Level-1 trigger rate for single and di-jet triggers, together with triggers including $\Delta\eta_{jj}$ and $\Delta\phi_{jj}$ cuts. These topological cuts on the di-jet trigger reduce the trigger rate dramatically, by almost an order of magnitude for a 100 GeV threshold. The corresponding generator level rates are also shown for comparison.

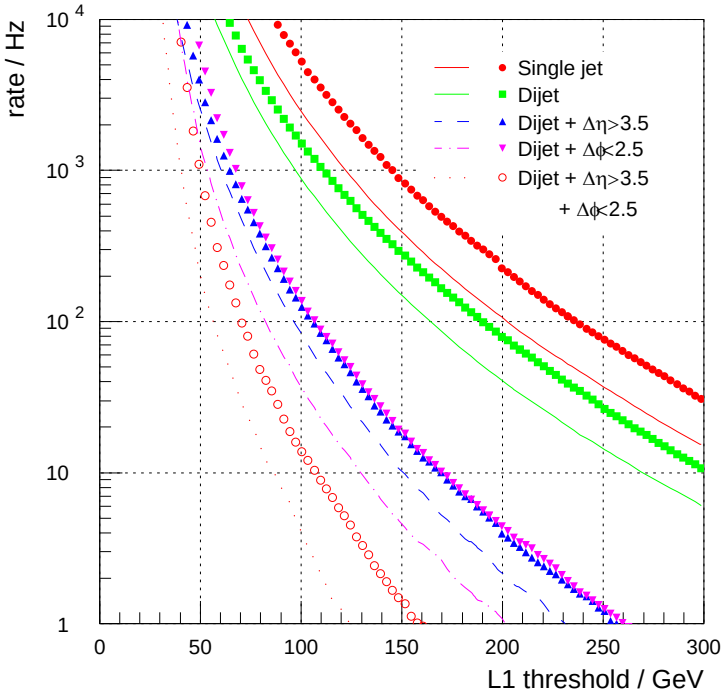

Figure 7.14 - Level-1 Trigger rate for various triggers, with and without topological cuts. The smooth curves show the equivalent generator level rate.

7.4.3 Trigger Efficiency

Figure 7.15 shows the Level-1 trigger rate plotted against the WBF event selection efficiency for single jet, di-jet, and three 'di-jet + topological cut' triggers. The efficiency is calculated with respect to events passing the tag quark cuts, and the rate is calculated from the binned minimum-bias samples. As can be seen, the single jet trigger provides the best efficiency for a given rate, except below ~800Hz, where the 'di-jet + $\Delta\eta$ + $\Delta\phi$' trigger provides marginally better efficiency, though this remains very poor.

The same triggers are plotted in Figure 7.16, but here the efficiency is calculated with respect to events passing the generator level $\Delta\eta$ cut in addition to the tag quark cuts. Here, the 'di-jet

+ Δη' trigger provides better performance than the previous plot. For the most part (up to ~3 kHz rate) it is as good as the single jet trigger. Also, the 'di-jet + Δη + Δϕ' trigger provides improved performance and is better than the single jet trigger up to a rate of ~1 kHz.

The selection of events passing the generator level Δϕ cut in addition to the tag quark cuts is shown in Figure 7.17. The 'di-jet + Δη + Δϕ' trigger might be expected to provide better performance here than it does. In fact, the 'di-jet + Δη' trigger provides almost as good performance.

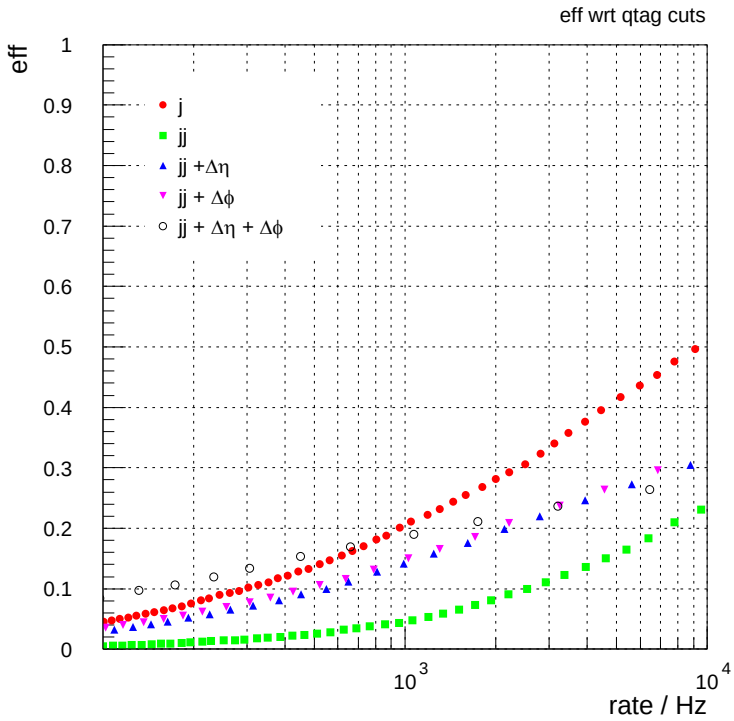

Figure 7.15 - Rate vs. Efficiency for several L1 triggers. Efficiency is calculated w.r.t the WBF sample after 'tag quark' cuts.

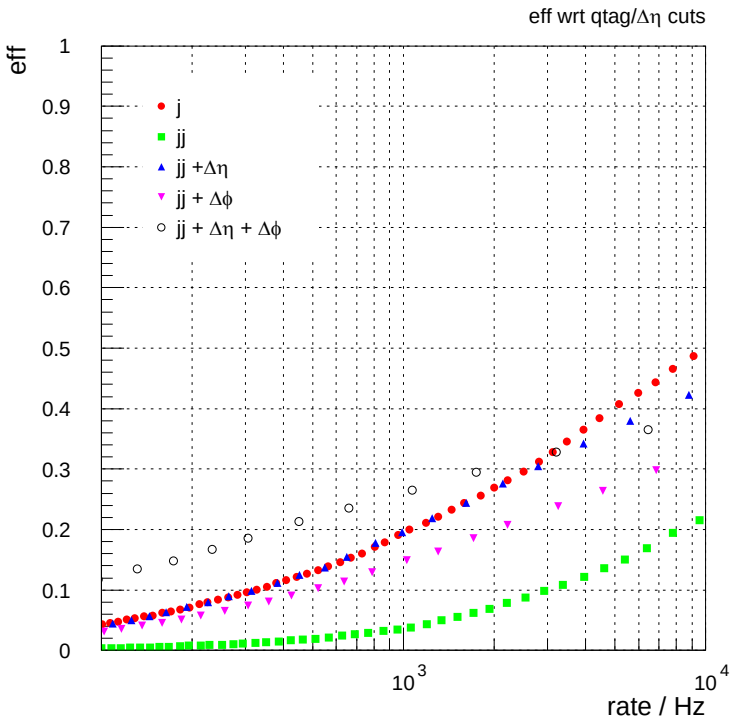

Figure 7.16 - Rate vs. Efficiency for L1 triggers. Efficiency is calculated w.r.t. the WBF sample after 'tag quark' and '$\Delta\eta$' cuts.

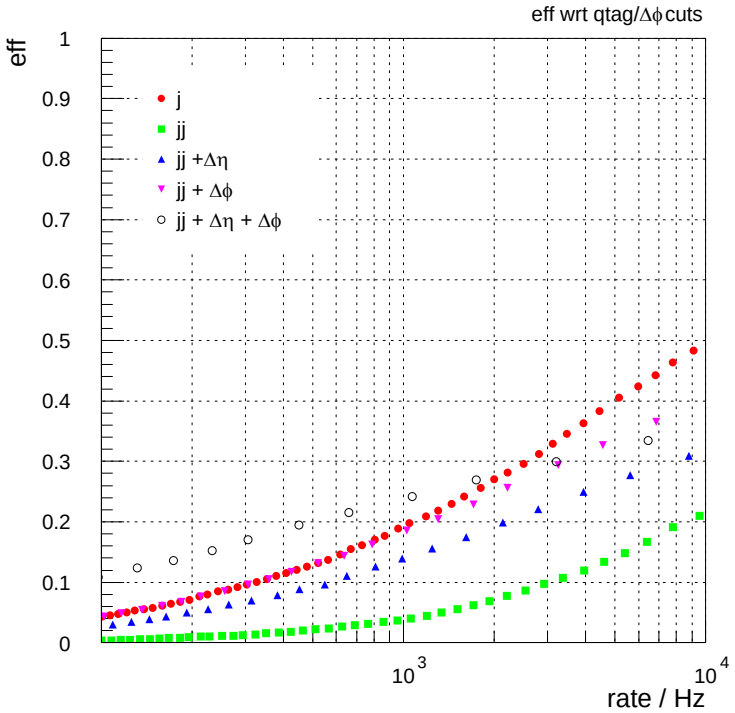

Figure 7.17 - Rate vs. Efficiency for L1 triggers. Efficiency is calculated w.r.t the WBF sample after 'tag quark' and '$\Delta\phi$' cuts.

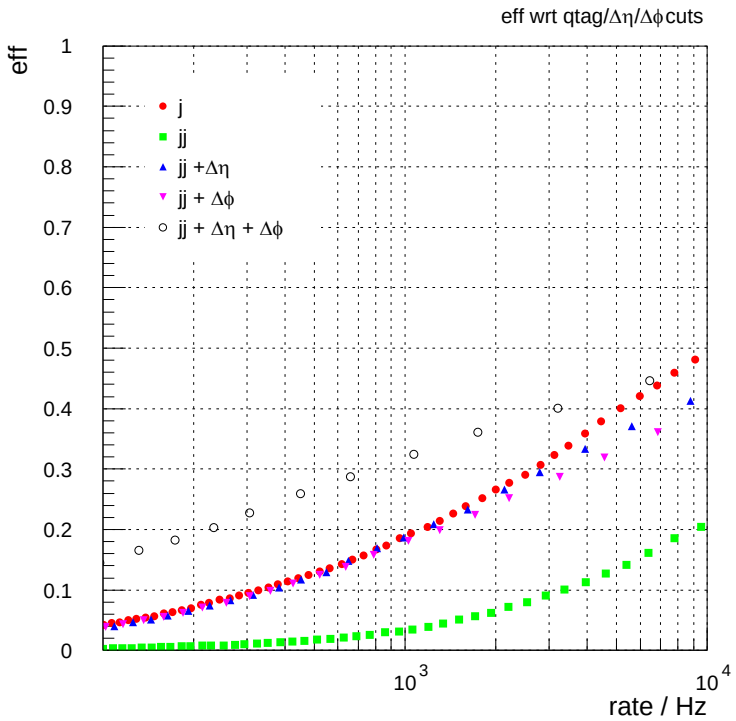

Figure 7.18 - Rate vs. Efficiency for L1 triggers. Efficiency is calculated w.r.t the WBF sample after 'tag quark', 'Δη' and 'Δϕ' cuts.

The Δφ and Δη triggers select tag quark events equally well (or badly), since they both essentially select events at the upper end of the Higgs E_t spectrum. It is therefore unsurprising that the Δφ + Δη trigger select such events better still.

The general efficiency remains extremely poor, however, when compared with the generator level results in Figure 7.13. This is the result of a combination of factors, including the tag jet identification efficiency and the Level-1 E_t, η and φ resolution. The generator level results showed the di-jet trigger with topological cuts to result in substantially greater efficiency than the single jet trigger, but this is not reflected in the Level-1 results. This is due to the difference in tag quark selection efficiency; the single jet trigger gives 93% efficiency compared to 67% for the di-jet trigger.

7.4.4 Extending the Tag Jet Identification

The di-jet + topological cut triggers clearly require correct identification of the tag jets to provide good efficiency. The triggers described so far have relied on correct selection of both tag jets by the 1st and 2nd Level-1 jets. This is only true for ~67% of all cases. In an attempt

to improve this efficiency, the next best combination of Level-1 jets was included. Now each trigger consists of an OR of the two dijet + topological cut conditions, i.e.

$$(E_t(J_2) > X + \Delta\eta(J_1J_2) > 3.5) \,\|\, (E_t(J_3) > X + \Delta\eta(J_1J_3) > 3.5) \tag{7.8}$$

$$(E_t(J_2) > X + \Delta\phi(J_1J_2) < 2.5) \,\|\, (E_t(J_3) > X + \Delta\phi(J_1J_3) < 2.5) \tag{7.9}$$

$$\begin{aligned}&(E_t(J_2) > X + \Delta\eta(J_1J_2) > 3.5 + \Delta\phi(J_1J_2) < 2.5) \,\|\\ &(E_t(J_3) > X + \Delta\eta(J_1J_2) > 3.5 + \Delta\phi(J_1J_3) < 2.5)\end{aligned} \tag{7.10}$$

Figures 7.19 - 7.22 show trigger efficiency as a function of rate for the four sets of generator level cuts, after inclusion of the extra tag jet pair. When compared with Figures 7.15 - 7.18, it is clear that the efficiency is either unaffected, or in most cases worsens, after including the extra jet pair. The reason for this is that the rate increases slightly for all E_t thresholds, but the efficiency only improves where the 1st and 3rd jets correctly select the tag jets. In general, the latter is only true at low energies, hence an improvement in efficiency should only be expected for low thresholds (and therefore high rate). The thresholds for the rate range we consider here (100 Hz - 10 kHz) are well above this regime, hence a decrease in efficiency is observed.

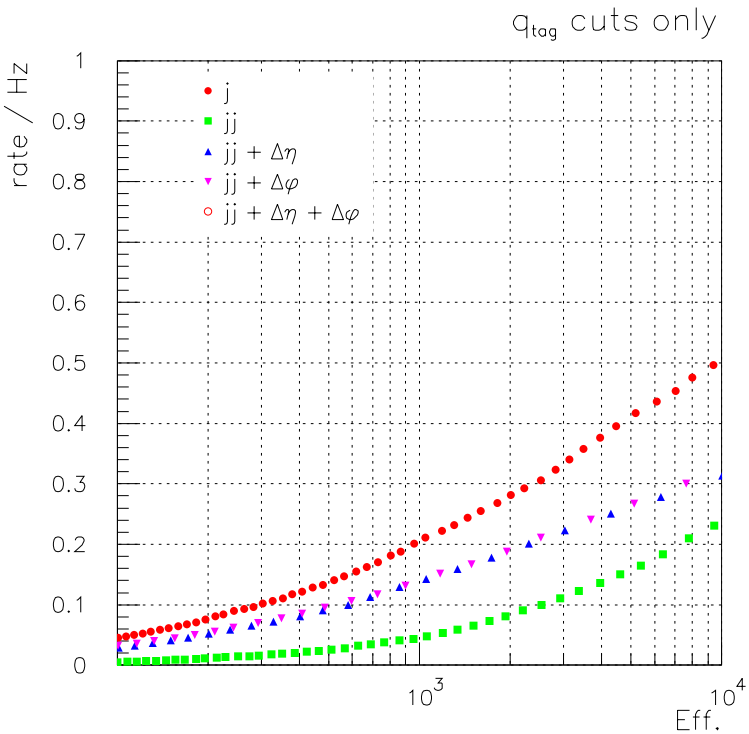

Figure 7.19 - Rate vs. efficiency (including extended Level-1 tag jet ID) after generator level 'tag quark' cuts

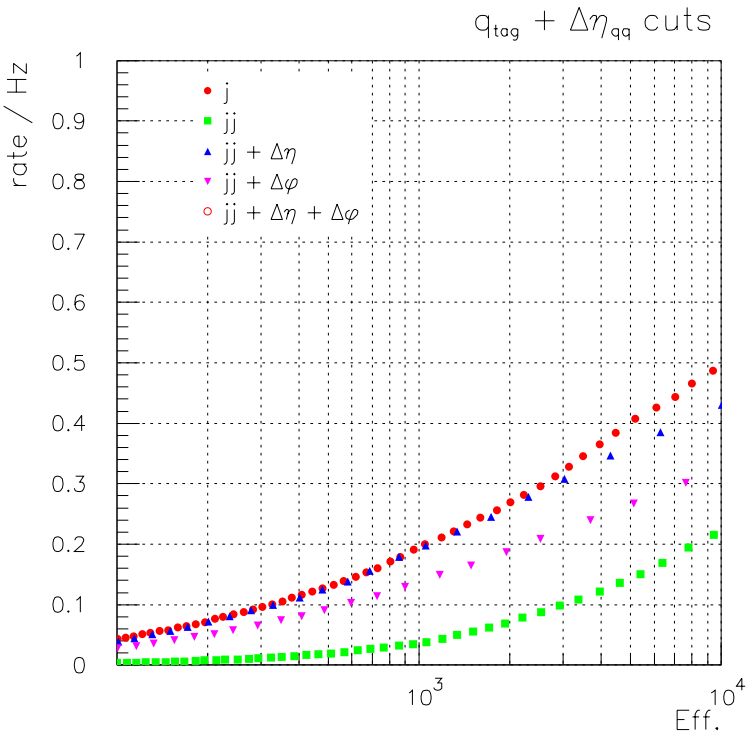

Figure 7.20 - Rate vs. efficiency (including extended Level-1 tag jet ID) after generator level 'tag quark' and 'Δη' cuts

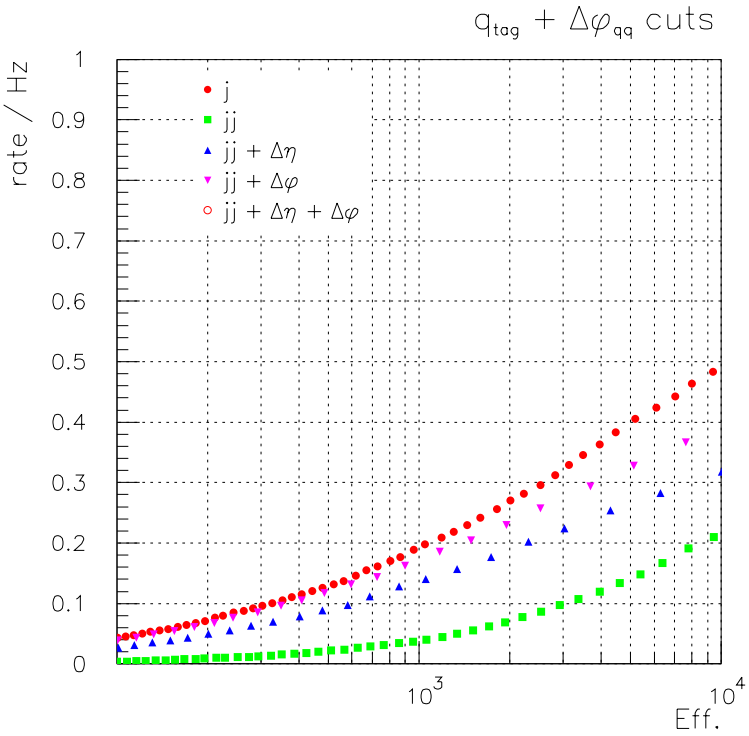

Figure 7.21 - Rate vs. efficiency (including extended Level-1 tag jet ID) after generator level 'tag quark' and 'Δϕ' cuts

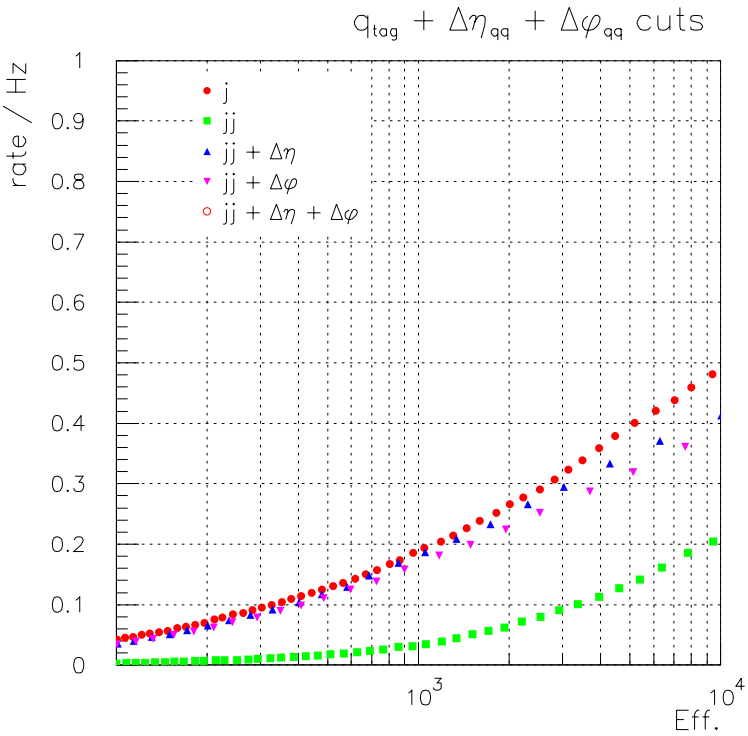

Figure 7.22 - Rate vs. efficiency (including extended Level-1 tag jet ID) after generator level 'tag quark', 'Δη' and 'Δϕ' cuts

The distribution of the $\Delta\eta(J_1J_3)$ and $\Delta\phi(J_1J_3)$ variables is shown in Figure 7.23. These plots confirm that these variables provide poor discrimination between minimum-bias events and the WBF signal.

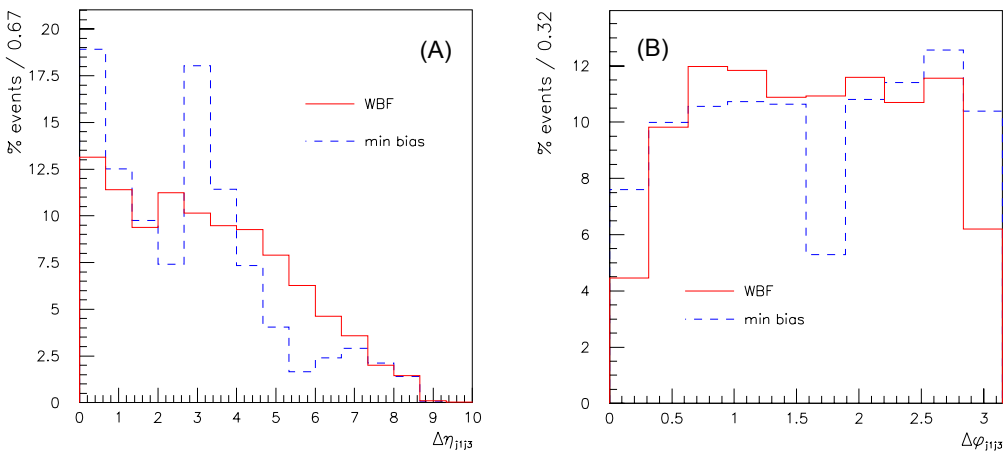

Figure 7.23 - Spacial separation (in η and ϕ) of 1st and 3rd Level-1 jets, for minimum bias and WBF samples.

In conclusion, then, the topological cuts do indeed improve the di-jet trigger efficiency for WBF events. However, because the tag-jet identification efficiency is substantially lower for di-jets than for the single jet trigger, the single jet trigger outperforms the topological di-jet trigger.

7.5 Invisible Higgs Triggers

In the last section, it was shown that the tag jets cannot be used alone to trigger on WBF processes. However, it is likely that triggers based on the combination of tag jets as well as the decay products of the Higgs will prove more performant than either the pure tag jet triggers, or the pure decay product triggers. In this section, the Level-1 missing E_t calculation is used in conjunction with tag jets to trigger on the WBF signal where the Higgs decays invisibly. The aim is to construct a trigger that selects the invisible Higgs signal with as high efficiency as possible, at as low a rate as possible

A variety of scenarios were outlined in §2.4, in which the Higgs may have a sizeable decay branching ratio to invisible states. Even a standard model Higgs may decay invisibly via ZZ* to neutrinos. For a light Higgs, however, the branching ratio for this process is extremely small, as the Z* is highly virtual. Weak boson fusion may present our best chance of detecting an invisibly decaying Higgs at the LHC. Clearly, the gluon-gluon fusion production mode offers no possibility of detection, and the other associative production modes (ttH, WH, ZH etc.) have lower cross-sections than WBF (see Figure 7.1).

This section will proceed as follows. First, the selection efficiency provided by the tag jet triggers described above is calculated for this signal. Higher efficiency is expected, since the offline analysis requires substantial transverse energy of the Higgs (and hence higher transverse energy of the tag jet). Next, the inclusion of a Level-1 threshold on missing energy is investigated, followed by asymmetric E_t thresholds on the di-jets. Finally, the full trigger efficiency is evaluated.

7.5.1 Event Selection

In the proposed ‘invisible Higgs’ offline analysis, initial event selection is made on the basis of the reconstructed jets and missing energy. The trigger efficiency is clearly best measured with respect to such cuts, as there is little point in achieving good efficiency for events that will be rejected immediately afterwards. However, since the offline reconstruction algorithms are not yet finalised, events are selected here on the basis of generator level quantities. The cuts given below are referred to later as the ‘Invisible Higgs’ cuts.

$$2 \text{ quarks}, E_t > 40 \text{ GeV}, |\eta| < 5.0 \quad (7.11)$$

$$E_t^{miss} > 100 \text{ GeV} \quad (7.12)$$

$$|\Delta\eta_{qq}| > 4.4 \quad (7.13)$$

$$|\Delta\phi_{qq}| < 1 \quad (7.14)$$

These cuts will be referred to later as 'invisible Higgs' cuts.

The missing energy cut is made on the quantity calculated after the application of a fiducial cut of $|\eta| < 5$ on the contributing particles. Figure 7.24 shows the quark and missing E_t spectra, calculated at generator level, after cuts 7.11 and 7.12. The tag quark spectrum is made substantially harder by the missing energy cut. The $\Delta\eta$ and $\Delta\phi$ distributions are also affected by the missing energy cut, as shown in Figure 7.26.

It is clear from the $\Delta\phi$ distribution after the harder cuts on tag quark E_t and E_t^{miss} that the $\Delta\phi$ trigger may provide greater efficiency than before. However, such triggers are not considered here because the offline analysis may be compromised. In order to estimate well the Wjj and Zjj backgrounds after cut 7.14, the background is measured using leptonic decays of the W and Z before the cut. Although the cut of $\Delta\phi < 2.5$ may not harm such estimates, further study would be needed in order to demonstrate this.

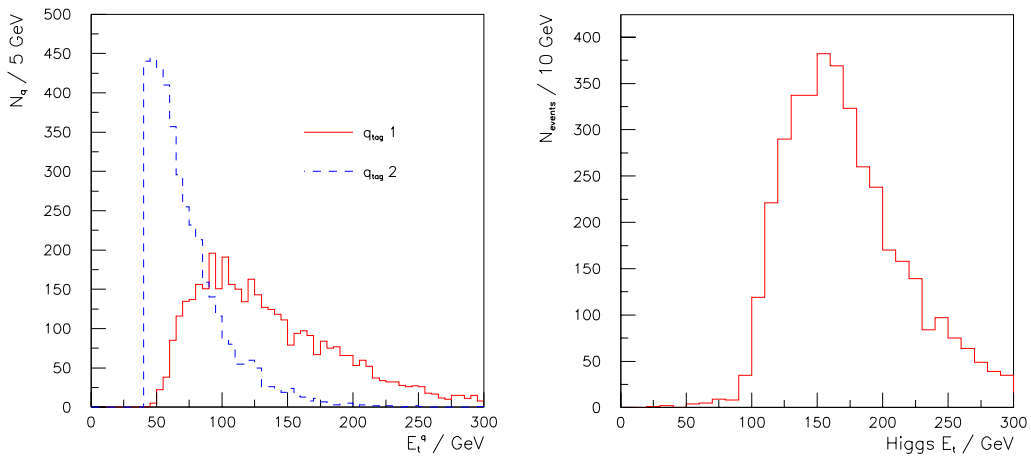

Figure 7.24 - Tag quark and Higgs E_t spectra of invisible Higgs signal after generator level cuts 7.11 and 7.12.

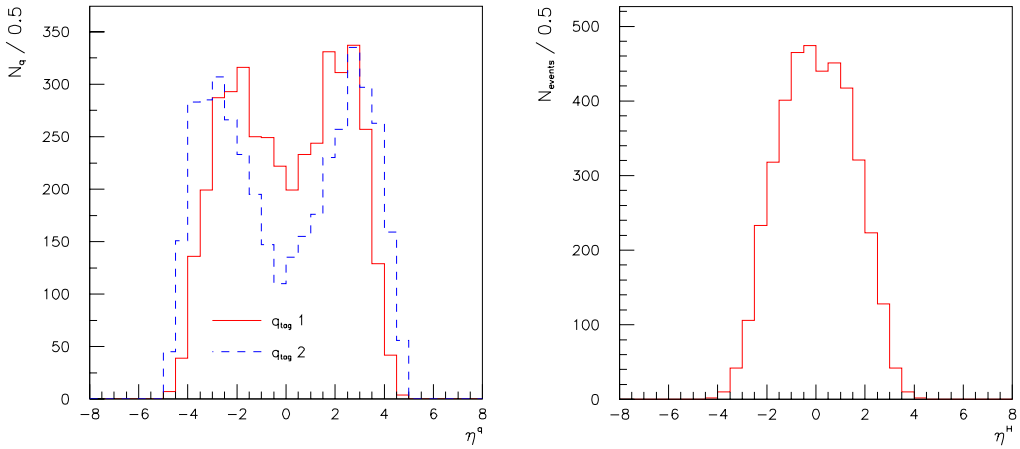

Figure 7.25 - Distributions of tag quarks and Higgs in pseudorapidity after invisible Higgs generator level cuts

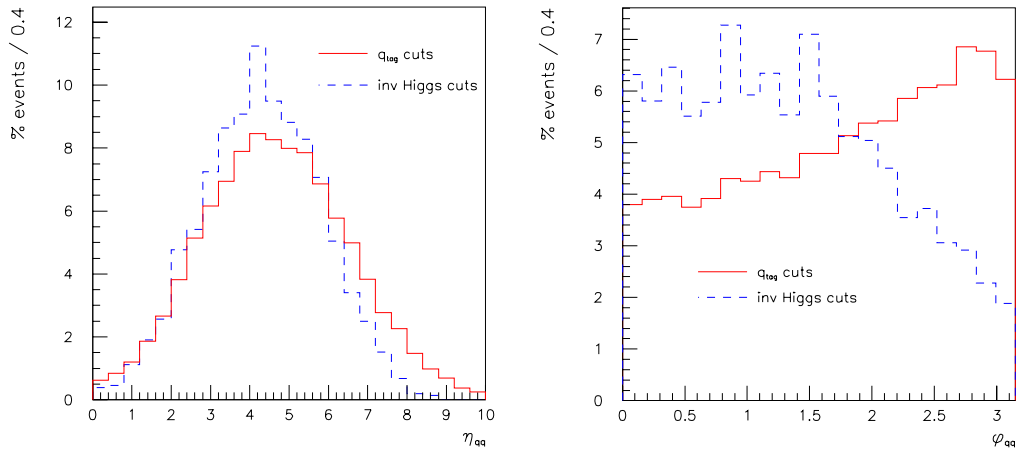

Figure 7.26 - Tag quark $\Delta\eta$ and $\Delta\phi$ distributions of the invisible Higgs signal, shown separately for tag quark and invisible Higgs generator level cuts.

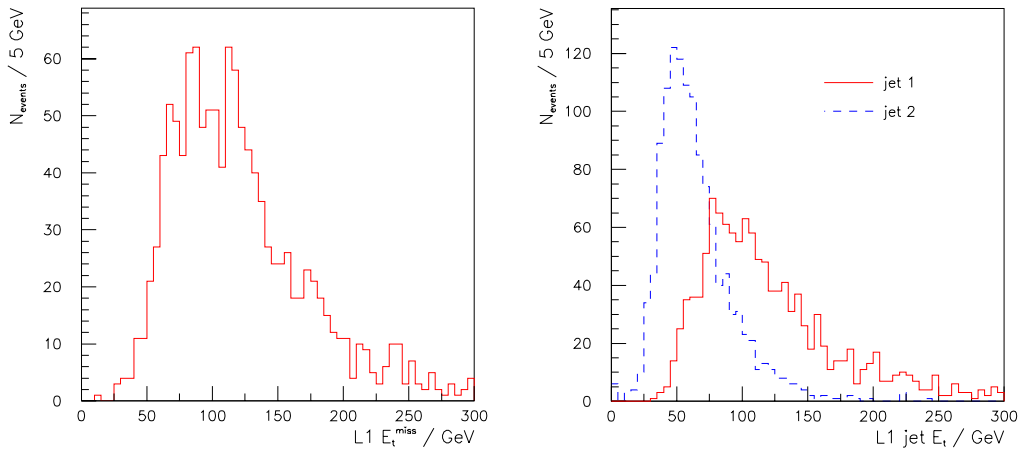

Figure 7.27 - Missing E_t and jet E_t distributions, measured by

Level-1, from the invisible Higgs signal after generator level cuts.

7.5.2 Performance of Simple Triggers

The efficiency of several simple triggers from Table 5.4 are plotted as a function of rate in Figure 7.28. As can be seen, the E_t^{miss} trigger provides substantially better performance than the jet based triggers. However, only a modest rate (10 Hz) is allotted to the missing E_t trigger. The efficiency of these triggers at the thresholds given in Table 5.4 are listed in Table 7.2.

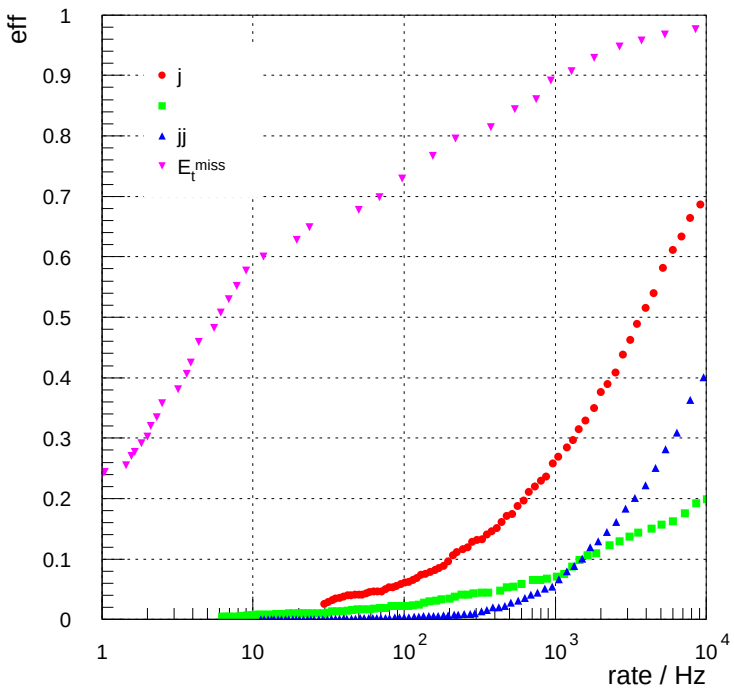

Figure 7.28 - Rate vs. efficiency for simple triggers from the Level-1 table.

Trigger	Thresh / GeV	Rate / kHz	Eff (ind)	Eff (cum)
j	110	3.2	47.2	47.2
jj	90	2.1	14.5	48.8
τ	85	4.9	15.5	53.7
E_t^{miss}	100	0.01	58.0	65.4

Table 7.2 - Invisible Higgs efficiency from simple Level-1 trigger table for low luminosity.

The baseline trigger performance is reasonably good, with a total efficiency of ~65%.

However, this can be improved upon by adding combined triggers that apply thresholds to both the missing energy, and to the tag jets. It is interesting to note that the 'di-jet + Δη' trigger now provides better efficiency, with respect to the invisible Higgs cuts, than the single jet trigger, as shown in Figure 7.29. This improvement is due to increased tag jet identification efficiency in the di-jet trigger, which is in turn due to the harder tag jet spectrum after the missing E_t cut. Despite this improvement, however, the 'di-jet + Δη' trigger still requires significant rate for useful efficiency (around half that required by the single jet trigger), and the combined triggers offer far greater performance, as will be shown in the next section.

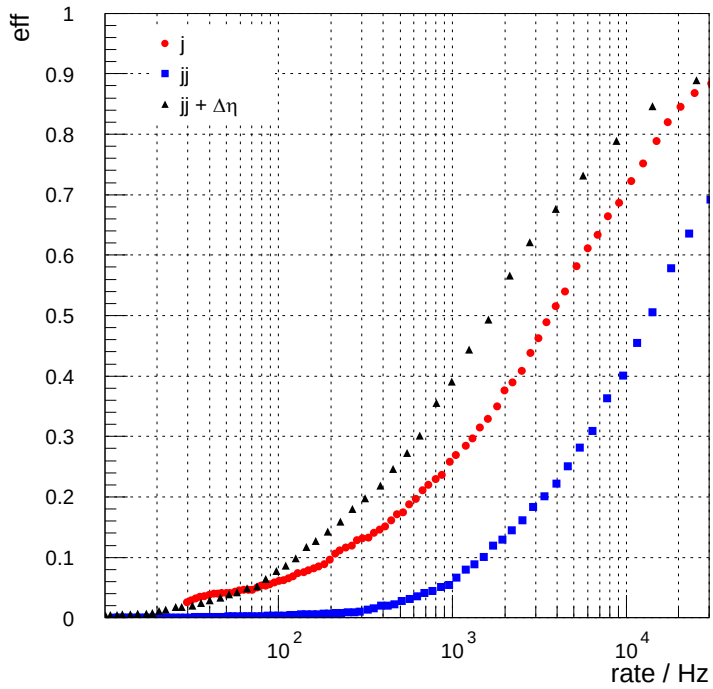

Figure 7.29 - Efficiency, with respect to invisible Higgs cuts, as a function of trigger rate for the single jet, di-jet and 'di-jet + Δη' triggers.

7.5.3 Tag jet + missing energy Triggers

By including a condition on the Higgs final state, we may be able to reduce the trigger rate sufficiently that the topological di-jet triggers become more useful. Here, a cut on the missing transverse energy is used to select invisible Higgs events. Figures 7.30 - 7.32 show rate and efficiency for jet triggers with the addition of a missing energy threshold. For each of the three thresholds (50, 60 & 70 GeV) the rate and efficiency of 'jet + E_t^{miss}', 'di-jet + E_t^{miss}' and 'di-jet + $\Delta\eta$ + E_t^{miss}' triggers are plotted. In each case, the 'di-jet + $\Delta\eta$ + E_t^{miss}' trigger offers substantial improvement over the other two triggers.

Note that as the E_t^{miss} threshold increases, steps begin to appear in the rate curve. This is believed to be due to a small number of unphysical events in the minimum bias sample used to calculate the rate. These events have a large energy deposit in a single ECAL crystal or HCAL cell, and are caused by a wrongly set parameter in the CMSIM program. This caused a large, unphysical energy deposit in a single calorimeter cell for a few events, and hence creates an unphysical missing E_t. The most extreme examples of such events were removed before calculating the rate plots in §5.5.2, but the number of more minor events meant that they could not be removed as easily. Since these events all contain more energy than they are supposed to, they serve to increase the rate rather than decrease it, and consequently any conclusion drawn from the rate plot will be more pessimistic than optimistic.

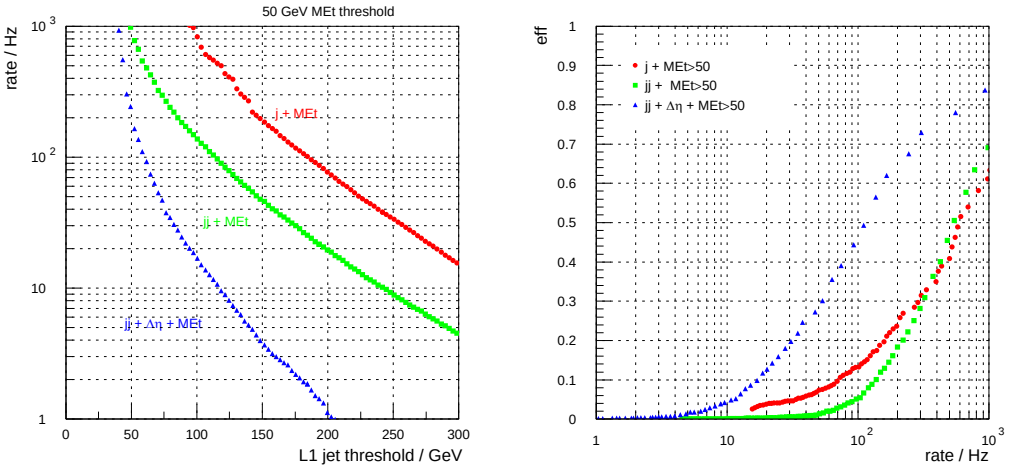

Figure 7.30- Rate and efficiency plots for jet triggers including a 50 GeV Etmiss threshold.

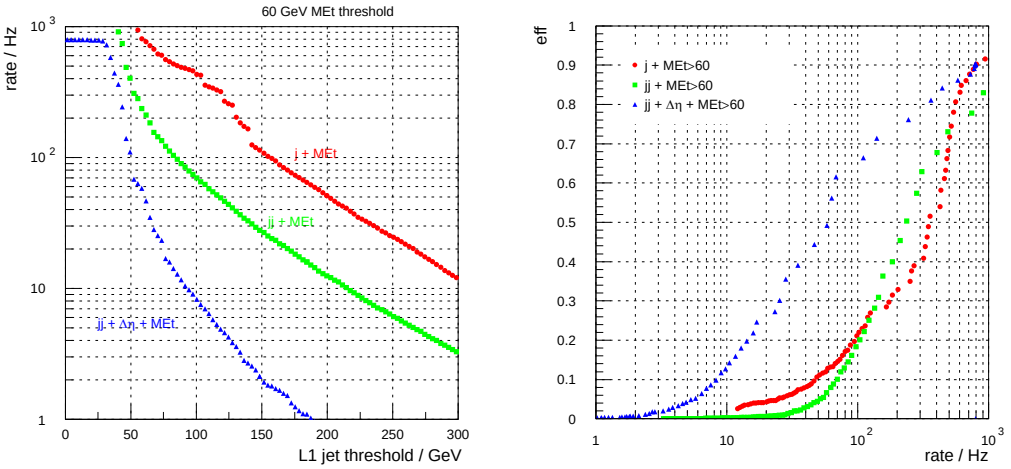

Figure 7.31- Rate and efficiency plots for jet triggers including a 60 GeV Etmiss threshold.

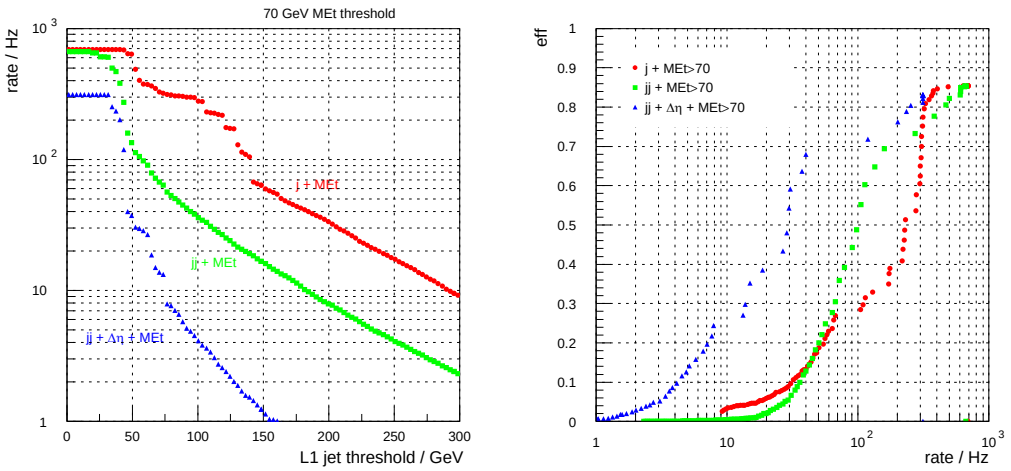

Figure 7.32 - Rate and efficiency plots for jet triggers including a 70 GeV Etmiss threshold.

The optimum missing energy threshold in the 'dijet + Δη + E_t^{miss}' trigger is determined from Figure 7.33. The 50 GeV trigger turns on too slowly, and the 70 GeV trigger plateaus too low at around 70% efficiency, so the 60 GeV cut is chosen.

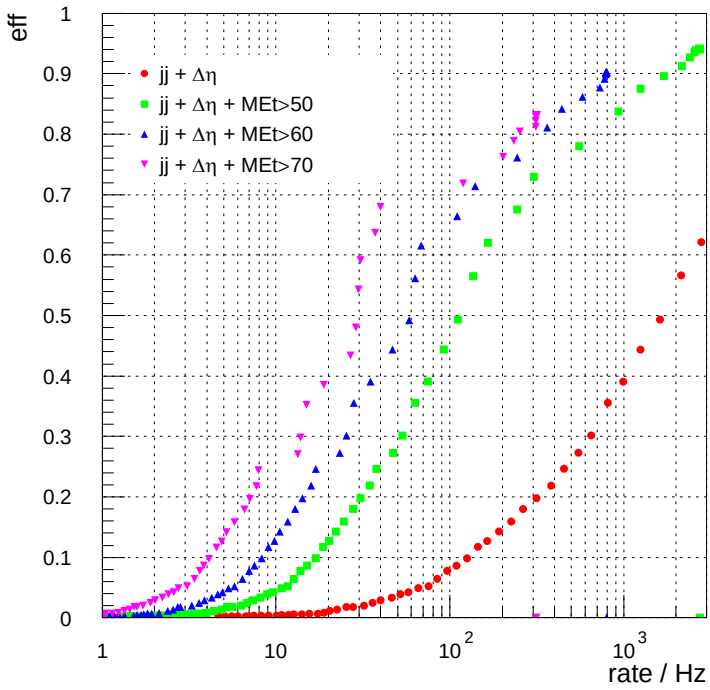

Figure 7.33 - Trigger efficiency as a function of rate for 'dijet + Δη + E_t^{miss}' triggers, with E_t^{miss} thresholds of 50, 60 & 70 GeV.

7.5.4 E_t threshold of 1st jet

The dijet threshold in the 'di-jet + Δη + E_t^{miss}' triggers is sufficiently low that it may be advantageous to include an additional threshold on the jet with highest E_t. The E_t distribution of this jet can be seen in Figure 7.24 to have a lower bound at around 50 GeV. The effect of including such a threshold is shown in Figure 7.34. Here, the invisible Higgs efficiency is plotted as a function of rate for 'di-jet + E_t^{miss}' triggers with and without both Δη and highest jet E_t cuts. A small improvement (up to 5%) in efficiency at the upper end of the trigger rate can be seen. Though this improvement is small, it is sensible to include such a cut if 'di-jet + E_t^{miss}' triggers with low di-jet thresholds are to be used.

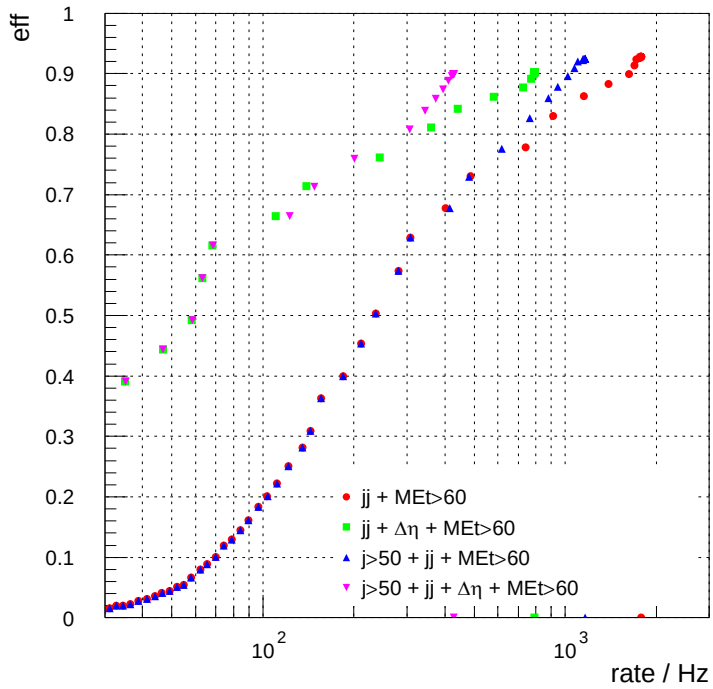

Figure 7.34 - Rate vs. efficiency for 'dijet + Δη' and 'dijet + Δη + Etmiss' triggers, with and without a 50 GeV threshold on the 1st jet.

A selection of 'jet + E_t^{miss}' triggers are listed in Table 7.3, with the individual and total rate and efficiency. The total rate and efficiency for each trigger are those obtained using the triugger in question in addition to the simple triggers from Table 5.4.

The E_t^{miss} threshold is set at 60 GeV for all triggers. The jet thresholds were then set by requiring ~90% total efficiency for invisible Higgs. As can be seen, the most perfomant trigger is the 'di-jet + Δη + E_t^{miss}' trigger with asymmetric di-jet thresholds, as this achieves the desired 90% efficiency at the lowest cost in terms of rate.

Trigger	Thresholds (GeV)	Rate (Hz)	Individ. Eff. (%)	Total Rate (Hz)	Total. Eff. (%)
j · E_t^{miss}	60, 60	800	88.9	7,350	89.1
jj · Δη · E_t^{miss}	33, 60	725	86.2	7,429	89.6
jj · Δη · E_t^{miss}	50, 30, 60	410	87.4	7,253	90.4

Table 7.3 - Invisible Higgs triggers. Total efficiency is calculated after the inclusion of j, jj, τ, ττ and E_t^{miss} triggers from Table 5.4.

7.6 Conclusion

The addition of topological cuts to the di-jet trigger have been studied, specifically cuts on the di-jet separation of $\Delta\eta_{jj} > 3.5$ and $\Delta\phi_{jj} < 2.5$. These triggers are found not to offer improved performance in selecting weak boson fusion events over the single jet trigger. The main reason for this is the tag jet identification efficiency at Level-1. The single jet trigger is found to correctly identify a tag jet with 93% efficiency, while the di-jet trigger successfully identifies both tag jets with 67% efficiency.

The 'di-jet + $\Delta\eta$' trigger may be of use, however, in diffractive physics studies where various measurements require two jets with substantial separation in pseudorapidity [63]. The thresholds required here are very low (>30 GeV), and will therefore require pre-scaling, but this trigger should provide a sample of greater purity than a simple pre-scaled di-jet or single jet trigger.

Several version of the 'jet(s) + missing E_t' trigger have been studied, where the jet requirements range from single jet, through di-jet to 'di-jet + $\Delta\eta_{jj} > 3.5$'. The performance of these triggers in selecting weak boson fusion events where the light Higgs decays invisibly has been assessed. In general, these combined triggers provide individual efficiency better than 80%, and when combined with the simple triggers from Table 5.4 the total efficiency is between 85% and 90%. The most performant trigger is the 'missing E_t + di-jet + $\Delta\eta_{jj} > 3.5$' trigger, using different E_t thresholds on the two jets. This trigger alone provides 87% efficiency with a Level-1 accept rate of 410 Hz, and the total efficiency, after including single jet, di-jet, single tau and E_t^{miss} triggers is 90%. However, the 'jet + E_t^{miss}' trigger provides almost as good efficiency at a slight cost in rate. Since this trigger will capture interesting physics of a more general nature, it makes a better candidate for inclusion in the Level-1 trigger table. However, inclusion of the 'di-jet + $\Delta\eta_{jj}$ + E_t^{miss}' trigger with topological cuts in conjunction with the 'jet + E_t^{miss}' trigger provides redundancy in the trigger table, and this allows for improved measurement of the trigger efficiency. It is therefore proposed that both trigger are included in the table. In this case it may be possible to increase the thresholds on the 'jet + E_t^{miss}' trigger without affecting the efficiency, and thereby saving some rate. Certainly, if after early data taking it the Higgs appears to have a sizeable invisible width, both triggers should be included in the table.

Chapter 8 - Conclusion

One of the many ingredients vital to successful collection of physics data with the CMS experiment is a fully functioning and reliable trigger system. In order that this data may then be properly analysed, the performance of the trigger system must be well understood. The studies described in this thesis form part of the work that must be carried out in order to achieve these two goals.

The Global Calorimeter Trigger relies entirely on programmable logic in order to fulfill its data processing requirements at a reasonable cost. The firmware development and hardware tests described in Chapter 4 demonstrate that current FPGA devices can be used for fast pipelined data processing as required by the GCT. Furthermore, the sort algorithm has been developed to the point where it is ready for implementation in the final system.

The jet trigger has been studied in depth, in order to ensure that it is capable of capturing the required physics data. These studies also constitute the first stage in quantising the performance of the final trigger system; as with any sub-detector, the trigger performance must be well understood in order that the real physics data can be properly analysed and interpreted. The jet trigger is found to perform well. The transverse energy resolution curve fits well a parameterisation using stochastic, constant and noise terms. The E_t resolution for jets with 100 GeV transverse energy is 28.6 GeV. The trigger turn-on is sharp, with the 95% efficient point being reached within 35 GeV of the threshold, for thresholds below 200 GeV. Finally, the single jet rate, at a threshold of 110 GeV is 3.2 kHz.

The missing transverse energy trigger also performs well. The resolution curve fits well a parameterisation as a function of total deposited transverse energy, including just a noise term. For 1 TeV total E_t, the missing E_t resolution is 29 GeV. Again, the trigger turn-on is good, with a 95% efficiency point being reached within ~90 GeV of threshold, for thresholds below 200 GeV. Though this is substantially more than the jet trigger, the missing E_t measurement is affected more seriously by pile-up energy, and this is as good as can be expected. A missing E_t threshold of 105 GeV gives an individual trigger rate of 10 Hz.

A study of trigger tower size in the forward calorimeters has been carried out, in order to quantise the difference between two proposed configurations. The performance of the forward jet trigger has been evaluated, again in terms of energy resolution, trigger turn-on and

rate. A small difference between the energy resolution of the two configurations is observed. When the resolution curves are fitted to a parameterisation, the stochastic and constant terms are found to be the same for both schemes, while the noise term reflects the different tower size as expected. Both segmentation schemes show good, and essentially identical, trigger turn-on curves. Some difference is seen between the trigger rate for the two schemes. This is explained in terms of the energy of fake jets. However, the difference in performance between the two schemes is sufficiently small that the final choice can be based on practical and financial criteria.

Finally, the use of topological cuts between the two jets of a di-jet trigger is investigated. Two signals are used to evaluate the trigger performance; a Higgs, produced via weak boson fusion, decaying to non-hadronic objects, and a Higgs produced in the same way that decays to purely invisible objects. Two cuts are evaluated, based on the separation of the tag jets in pseudorapidity and azimuthal angle. For the general case of WBF production, neither cut is found to outperform the single jet trigger. However, in specific case of an invisibly decaying Higgs, the cut in $\Delta\eta$ is found to be of use. The best trigger for selecting this channel is found to be a combined di-jet and missing energy trigger, incorporating the cut on jet separation in pseudorapidity.

References

The following abbreviations have been used below:

- LHWG - LEP Higgs Working Group
- LSWG - LEP SUSY Working Group
- ADLO - ALEPH, DELPHI, L3, OPAL

Chapter 2 - Theoretical Background

[1]P. Higgs, Phys. Lett. **12** (1964) p132

[2]G. t'Hooft, Nucl Phys **B35** (1971) p167

[3]LEP & SLD Electroweak Working Groups (presented by G. Myatt), *Electroweak Results and Fits to the Standard Model*, proceedings of XXXVIIth Rencontres de Moriond (March 2002)

[4]ADLO collaborations, *Search for the Standard Model Higgs Boson at LEP*, CERN-EP/**2001-055** (July 2001)

[5]Super-Kamiokande collaboration (Y. Fukuda et. al.), *Evidence for Oscillation of Atmospheric Neutrinos*, Phys Rev Lett **81** (1998) p1562

[6]SNO collaboration (Q.R. Ahmad et. al.), *Direct Evidence for Neutrino Flavor Transformation from Neutral Current Interactions in the Sudbury Neutrino Observatory*, Phys Rev Lett **89** (2002) 011301

[7]G. Gelmini & M. Roncadelli, Phys. Lett. **B99** (1981) p411

[8]K. Olive, *Introduction to Supersymmetry: Astrophysical and Phenomenological Constraints*, 'Les Houches 1999 : The Primordial Universe' (1999) p221

[9]J. Ellis, S. Kelley & D.V. Nanopoulos, *Precision LEP Data, Supersymmetric GUTS and String Unification*, Phys Lett **B249** (1990) p441

[10]D.E. Groom et al., The European Physical Journal **C15** (2000) p1 and 2001 off-year partial update for the 2002 edition available on the PDG WWW pages (URL: http://pdg.lbl.gov/)

[11]S. Weinberg, *The Quantum Theory of Fields vol. III - Supersymmetry*, Cambridge Univ. Press (2000)

[12]LSWG & ADLO collaborations, *Combined LEP Selectron/Smoun/Stau results 183-208 GeV*, LEP SUSYWG/**02-01.1** (2002)

[13]LSWG & ADLO collaborations, *Combined LEP Stop and Sbottom Results 183-208 GeV,* LEP SUSYWG/02-02.1

[14]LSWG & ADLO collaborations, *Combined LEP Chargino Results up to 208 GeV,* LEP SUSYWG/**02-04.1**

[15]LHWG & ADLO collaborations, *Searches for the Neutral Higgs Bosons of the MSSM, Preliminary Combined Results using LEP 2 Data Collected up to 209 GeV,* LHWG Note **2001-04** (July 2001)

[16]LHWG & ADLO collaborations, *Searches for Charged Higgs Bosons, Preliminary Combined Results using LEP 2 Data Collected up to 209 GeV,* LHWG Note **2001-05** (July 2001)

[17]LSWG & ADLO collaborations, *Combined Lower Bound on the Neutralino Mass in a Constrained MSSM*, LEP SUSYWG/**02-07.1**

[18]E281 muon (g-2) collaboration, *Measurement of the Positive Muon Anomalous Magnetic Moment to 0.7ppm*, **hep-ex/0208001**

[19]S.P. Martin, J.D. Wells, *Motivation and detection of an Invisibly Decaying Higgs boson at the Fermilab Tevatron*, Phys Rev D **60** (1999) 035006

[20]A. Djouadi, P. Janot, J. Kalinowski, P.M. Zerwas, *SUSY Decays of Higgs Particles*, Phys Lett **B376** (1996) p220

[21]Y. Chikasige, R. Mohaptra, R. Peccei, Phys. Lett. **B98** (1981) p265

[22]N. Arkani-Hamed, S. Dimopoulos, G. Dvali, *The Hierarchy Problem and New Dimensions at a Millimetre*, Phys Lett **B249** (1998) p263

[23]G. F. Giudice, R. Rattazzi, J.D. Wells, *Graviscalars from Higher-Dimensional Metrics and Curvature-Higgs Mixing*, CERN-TH **2000-051** (February 2000)

[24]M. Maltoni, V.A. Novikov, L.B. Okun, A.N. Rozanov, M.I. Vysotsky, *Extra Quark-Lepton Generations and Precision Measurements,* Phys Lett **B476** (2000) p107

[25]V.A. Khoze, *Comment on an invisible Higgs boson and a 50 GeV neutrino*, **hep-ph 0105069** (May 2001)

[26]LEP Higgs Working Group, *Searches for Invisible Higgs bosons: Preliminary Combined Results Using LEP Data Collected at Energies up to 209 GeV,* **hep-ex/0107032** (2001)

Chapter 3 - The CMS Detector

[27]The CMS Collaboration, *CMS : Technical Proposal*, CERN/LHCC **94-38** (Dec 1994)

[28]The CMS Collaboration, *CMS : Tracker Technical Design Report*, CERN/LHCC **98-6** (Apr 1998)

[29]The CMS Collaboration, *Addendum to the CMS Tracker TDR*, CERN/LHCC **2000-16** (Feb 2000)

[30]The CMS Collaboration, *CMS : Electromagnetic Calorimeter Technical Design Report*, CERN/LHCC **97-33** (Der 1997)

[31]P. Depasse et al., *Analysis of August 1999 beam tests of a PbWO4 crystal matrix*, CMS Note **2000/009**

[32]K. Lassila-Perini, *Reconstruction of Higgs to γγ in CMS*, CMS CR **1997/006**

[33]The CMS Collaboration, *CMS : Hadronic Calorimeter Technical Design Report*, CERN/LHCC **97-31** (June 1997)

[34]D. Green et al., *Dijet Mass Resolution of the CMS Calorimeter*, CMS Note 1998/026

[35]S. Abdullin, *Squark and gluino searches in lepton + jets + E_t^{miss} final states*, CMS TN **1996/095**

[36]P. Gorodetzky et al., *Quartz Fibre Calorimetry*, Nucl. Instr. and Meth. in Phys. Res. **A361** (1995) p161

[37]N. Akchurin et al., *Beam test results from a fine-sampling quartz fibre calorimeter for electron, photon and hadron detection*, Nucl. Instr. and Meth. in Phys. Res. **A399** (1997) p202

[38]N. Akchurin et al., *Test beam results of CMS quartz fibre calorimeter prototype and simulation of response to high energy hadron jets*, Nucl. Instr. and Meth. in Phys. Res. **A409** (1998) p593

[39]The CMS Collaboration, *CMS : Muon Technical Design Report*, CERN/LHCC **97-32** (Dec 1997)

[40]The CMS Collaboration, *CMS : Level-1 Trigger Technical Design Report*, CERN/LHCC **2000-038** (Dec 2000)

[41]S. Abdullin & S. Eno, *On Improvement of Level-1 Trigger E_t^{miss} Measurement*, CMS IN **2000/060** (Dec 2000)

Chapter 4 - An FPGA Sort Processor

[42]G. Heath, D. Newbold, C. Purves, U. Schäfer, *CMS Global Calorimeter Trigger Conceptual Design*, CMS TN **95/049** (1995)

[43]CMS Collaboration, *Level-1 Trigger Technical Design Report*, CERN/LHC **2000-038** (Dec 2000)

[44]U. Schäfer, private communication (1996)

Chapter 5 - Performance of the Level-1 Jet and Missing Energy Triggers

[45]S. Abdullin, *Search for SUSY in (leptons +) jets + E_t^{miss} final states*, CMS Note **1998/073** (Nov 1998)

[46]P. Chumney, S. Dasu & W. Smith, *CMS Level-1 Regional Calorimeter Trigger Simulation Results*, CMS IN **2000/074** (Dec 2001)

[47]P. Chmuney, S. Dasu, W. Smith, *Updated Regional Calorimeter Trigger Baseline Thresholds and Rates for a LHC luminosity of* $2 \times 10^{33}\ cm^{-2}s^{-1}$, CMS IN **2002/019** (Mar 2002)

[48]T. Sjöstrand et. al., *HEP Event Generation with PYTHIA 6.1*, Computer Phys.Commun. **135** (2001) p238

[49]CERN Application Software Group, *GEANT Detector Description and Simulation Tool*, CERN Long Writeup **W5013**

[50]C. Charlot et. al., *CMSIM - CMANA: CMS Simulation Facilities*, CMS TN **1993/63** (Oct 1993)

[51]D. Stickland et. al., *CMS Reconstruction Software: The ORCA Project*, CMS IN **1999/035** (Oct 1999)

[52]J. P. Wellisch, *An Object Oriented Jet Finder Library*, CMS Note **1999/034** (May 1999)

[53]P. Hidas & S. Kunori, *Level2 Missing Transverse Energy Trigger Studies at Low Luminosity*, CMS IN 2000/draft

[54]S.Abdullin et. al., *Energy Corrections for QCD Jets*, CMS IN **2001/001** (Jan 2001)

[55]A. Krokhotine, minutes of the JetMet PRS group meeting 11 Oct 2001, http://computing.fnal.gov/cms/jpg/minutes/oct112001/minutes.html

[56]P. Hidas, S. Abdullin & S. Eno, *Missing Transverse Energy Studies with ORCA4 on Minimum Bias and Signal Samples*, CMS Note **2001/005** (Sep 2000)

[57]A. Nikitenko, S. Kunori, R. Kinnunen, *Missing Transverse Energy Measurement with Jet Energy Corrections*, CMS Note **2001/040** (July 2001)

[58]J. Brason & E. Trepagnier, *Weighting bunch crossings to calculate rate in the presence of pile-up*, CMS IN **2000/038** (Aug 2000)

[59]H. Sakulin, *Global Muon Trigger Update*, presented at CMS Level-1 Trigger Meeting, (4 Dec 2001) (currently available at http://wwwhephy.oeaw.ac.at/p3w/cms/trigger/globalMuonTrigger/trans/GMT-CMSWeek04Dec2001.pdf)

Chapter 6 - Trigger Tower Size in the Very Forward Calorimeters

[60]A. Ferrando *et. al.*, *Detection and Reconstruction of Tagging Jets in (η, ϕ) Segmented Quartz Fibre Forward Calorimeters for CMS*, CMS Note **1998/015** (Feb 1998)

Chapter 7 - Topological Di-jet Triggers

[61]O. Eboli and D. Zeppenfeld, *Observing an invisible Higgs boson*, Phys. Lett. **B495** (2000) p147

[62]M. Spira, *QCD Effects in Higgs Physics*, CERN-TH **91-68** (1997)

[63]A. de Roeck, personal communication (2002)

Appendix A - Diagram of the CMS Detector

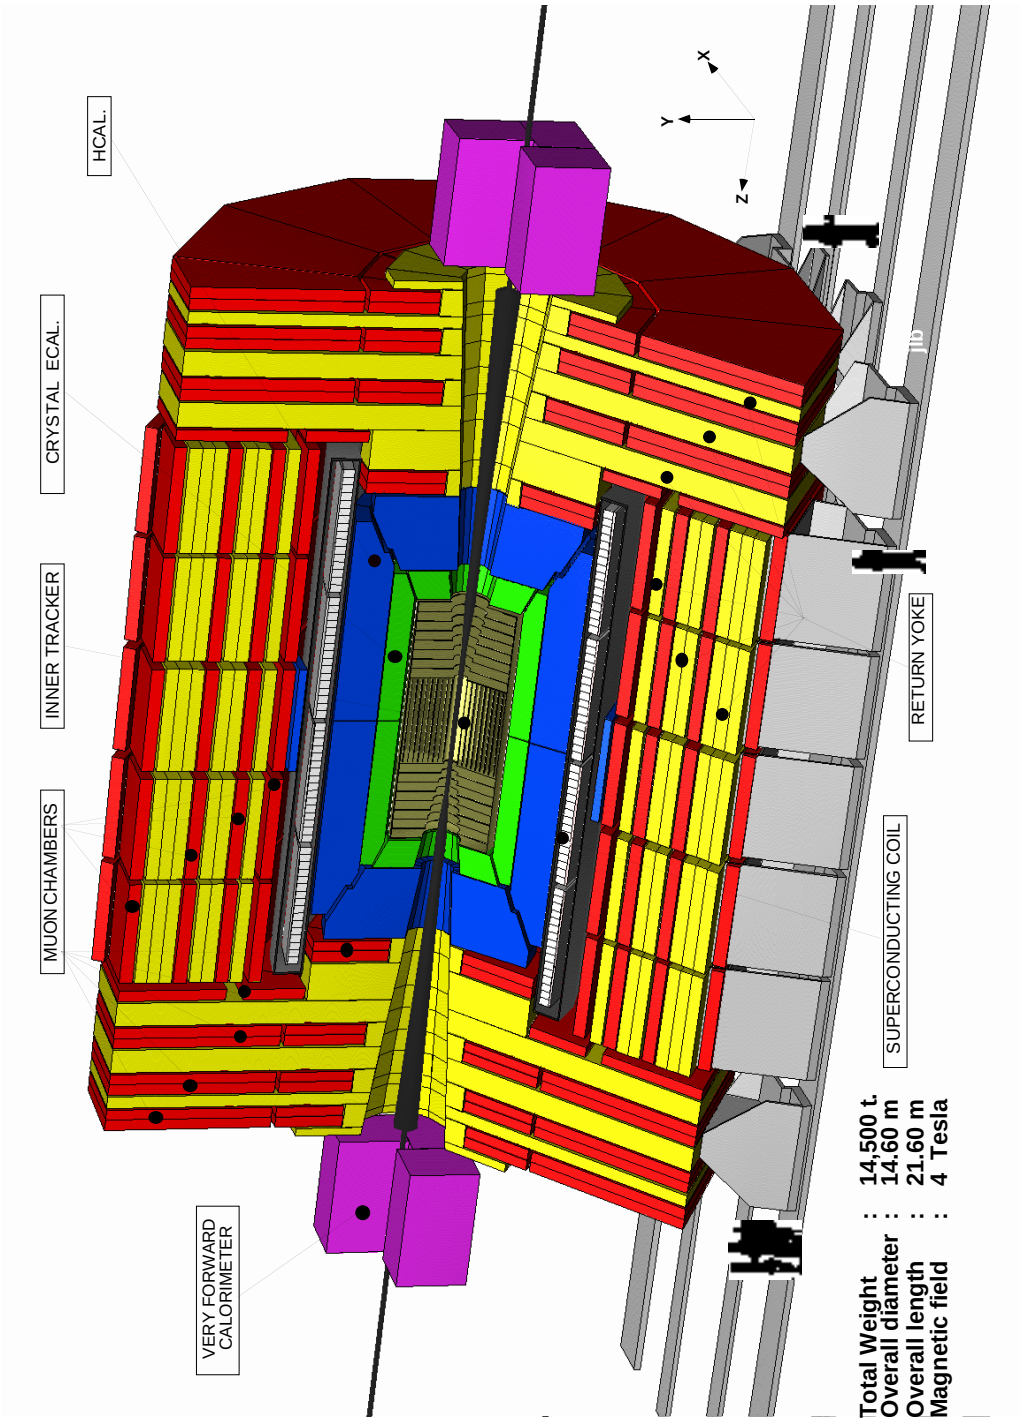

Figure A.1 - Cut away view of the CMS detector

Appendix B - Corrections to Level-1 Forward Jet Energy at High Luminosity

Corrections to the Level-1 jet transverse energy scale are given here for two different trigger tower segmentations of the very forward calorimeters. The two schemes correspond to trigger towers made up of 3×2 readout towers in (η, ϕ) and 2×2 readout towers in (η, ϕ). The transverse energy corrections are calculated by fitting a quadratic curve to the Level-1 E_t value as a function of the true (generator level) E_t value. i.e.

$$E_t^{L1} = a(E_t^{true})^2 + b(E_t^{true}) + c \tag{B.1}$$

The parameters a, b and c are found from the fit results and the reverse transform is applied to obtain the corrected E_t from the measured Level-1 E_t :

$$E_t^{corr} = \frac{\sqrt{b^2 - 4a(c - E_t^{L1})} - b}{2a} \tag{B.2}$$

Fits are performed for jets in each trigger region in η, as shown in Figure B.1 and Figure B.2 (for the 3×2 and 2×2 schemes respectively). The parameters resulting from each fit are also given in Table B.1 and Table B.2.

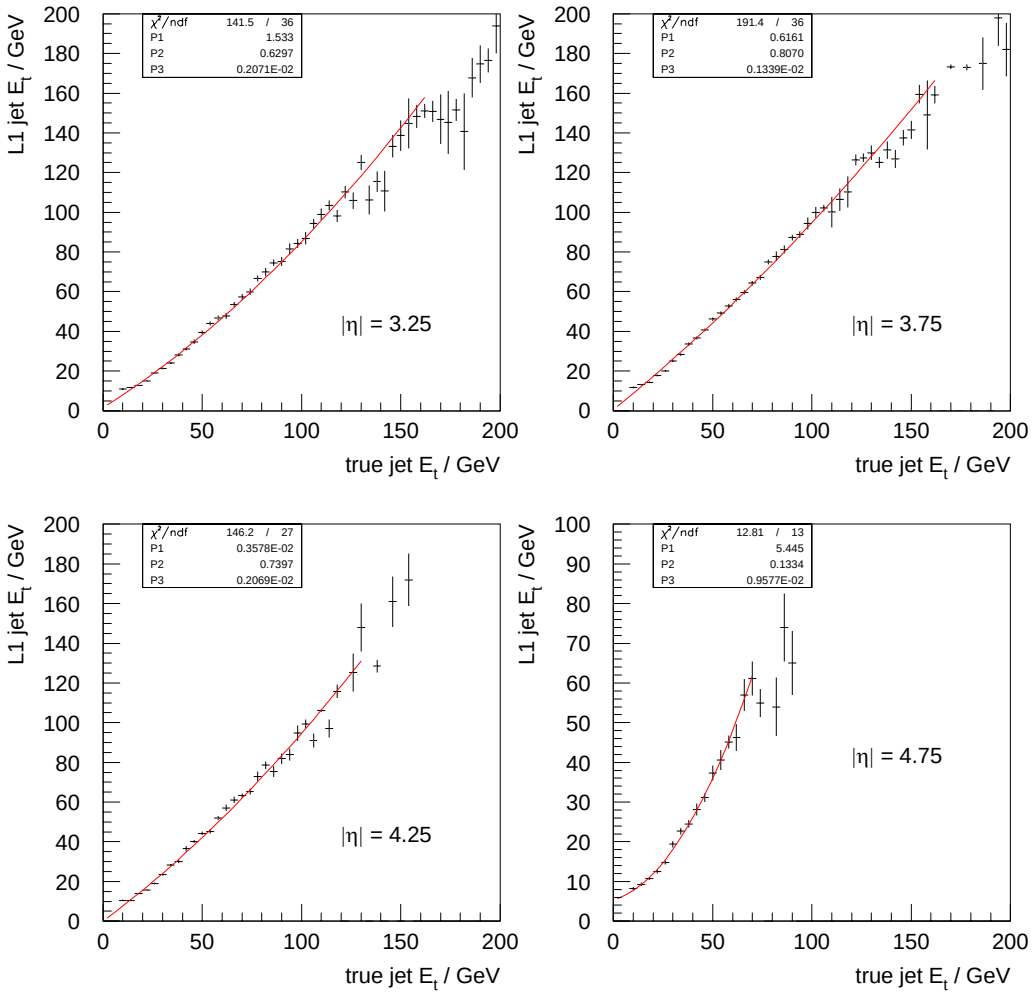

Figure B.1 - L1 forward jet E_t as a function of true jet E_t for each band of η in the 3×2 segmentation scheme

$\lvert\eta_{region}\rvert$	a / 10^{-3} GeV^{-1}	b / 10^{-1}	c / GeV
3.25	2.1 ± 0.2	6.3 ± 0.2	1.5 ± 0.4
3.75	1.3 ± 0.1	8.1 ± 0.2	0.6 ± 0.3
4.25	2.1 ± 0.2	7.4 ± 0.2	0.0 ± 0.4
4.75	9.6 ± 0.1	1.3 ± 0.7	5.4 ± 0.9

Table B.1 - Fit results for the 3×2 scheme.

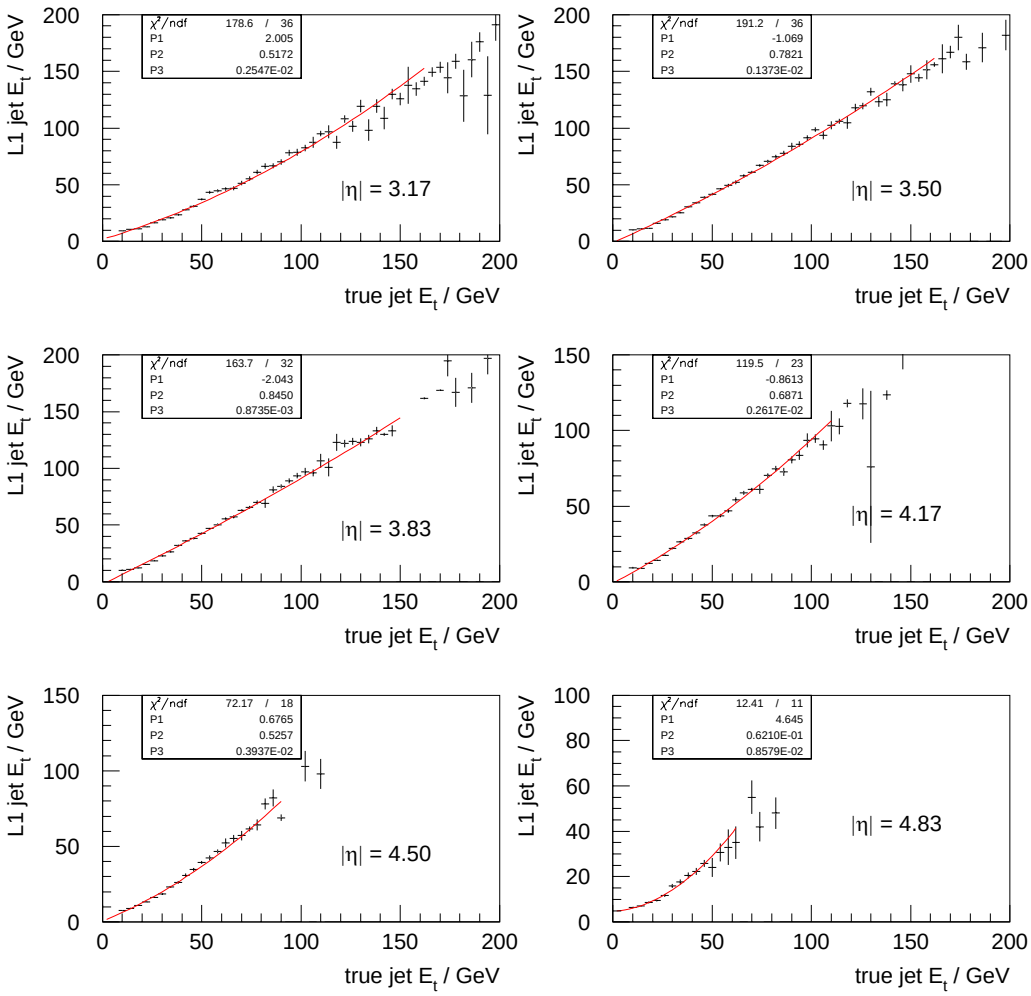

Figure B.2 - L1 forward jet E_t as a function of true jet E_t for each band of η in the 2×2 segmentation scheme

$\lvert\eta_{region}\rvert$	a / 10^{-3} GeV^{-1}	b / 10^{-1}	c / GeV
3.17	2.5 ± 0.2	5.2 ± 0.2	2.0 ± 0.3
3.50	1.3 ± 0.1	7.8 ± 0.2	-1.1 ± 0.3
3.83	0.9 ± 0.1	8.4 ± 0.2	-2.0 ± 0.4
4.17	2.6 ± 0.3	6.9 ± 0.3	-0.9 ± 0.4
4.50	3.9 ± 0.5	5.3 ± 0.4	0.7 ± 0.6
4.83	8.6 ± 1.6	0.6 ± 0.8	4.6 ± 0.9

Table B.2 - Fit results for the 2×2 scheme.